On the orientation of hypergraphs

by

Andres J. Ruiz-Vargas

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Combinatorics & Optimization

Waterloo, Ontario, Canada, 2010

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

This is an expository thesis. In this thesis we study out-orientations of hypergraphs, where every hyperarc has one tail vertex. We study hypergraphs that admit out-orientations covering supermodular-type connectivity requirements. For this, we follow a paper of Frank.

We also study the Steiner rooted orientation problem. Given a hypergraph and a subset of vertices $S \subseteq V$, the goal is to give necessary and sufficient conditions for an orientation such that the connectivity between a root vertex and each vertex of S is at least k, for a positive integer k. We follow a paper by Kiraly and Lau, where they prove that every $2k$-hyperedge connected hypergraph has such an orientation.

Acknowledgements

First of all, I would like to thank my advisor Professor Joseph Cheriyan for his support and guidance. His revisions helped to shape this thesis, as well as my understanding of the english language.

I would like to show my gratitude to the readers, Professor Bruce Richter and Professor William H. Cunningham, for taking the time to read this thesis and for their helpful suggestions.

I am grateful to all the people around me, for its them who give to this thesis a non mathematical importance in my life.

To my parents.

Table of Contents

Chapter 1

Introduction

1.1 Overview

This is an expository thesis. In this thesis we study out-orientations of hypergraphs, where every hyperarc has one tail vertex.

In Chapter 2, we study hypergraphs that admit out-orientations covering supermodular-type connectivity requirements. We follow a paper of Frank [3], where a result for graphs is proved. We prove such result for hypergraphs.

In Chapter 3, we study the Steiner rooted orientation problem. Given a hypergraph and a subset of vertices $S \subseteq V$, the goal is to give necessary and sufficient conditions for an orientation such that the connectivity between a root vertex and each vertex of S is at least k, for a positive integer k. We follow a paper by Király and Lau [7], where they prove that every $2k$-hyperedge connected hypergraph has such an orientation.

Before every major step in Chapter 2 and Chapter 3, we present a high-level explanation. Moreover, the thesis has detailed arguments for the last parts of Chapters 2 and 3, filling in some of the details missing from the original papers.

Throughout this thesis, k will denote a nonnegative integer. Throughout, we refer to a mathematical item numbered i (by latex) as Equation i; even if the item is not an equation. The remainder of Chapter 1 is devoted to definitions and results on hypergraphs. We also present some key results on orientation of hypergraphs.

1.2 Graphs and hypergraphs

In this section, we define graphs and hypergraphs.

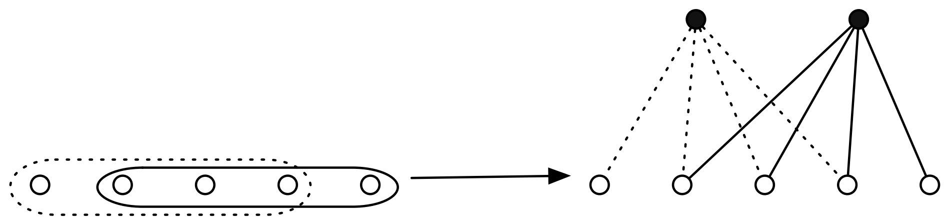

Figure 1.1: On the left, a hypergraph H with two hyperedges. On the right, its bipartite representation $B(H)$.

Definition 1. *A graph is a pair $G = (V, E)$ such that V is a set, and E is a multiset whose members are 2-element subsets of V.*

We only consider loopless graphs, as in the definition.

In this thesis, we are mainly interested in hypergraphs, a generalization of graphs, where instead of edges we have hyperedges. Every hyperedge is a subset of the set of vertices. That is, we drop the assumption that an edge has size 2. More formally,

Definition 2. *A hypergraph is a pair $H = (V, E)$ such that V is a set, and E is a multiset whose members are nonempty subsets of V.*

Hence, a hyperedge $e \in E$ is a nonempty subset of V of any size. Although the definition allows hyperedges of size 1, we assume that there are none. We let $V(H)$ denotes the set of vertices of H and $E(H)$ denotes the multiset of hyperedges of H.

For an easier way to picture a hypergraph, we will construct a graph that captures the essential properties of the hypergraph.

Given a hypergraph $H = (V, E)$, we say that $B(H) = (V_B, E_B)$ is the *bipartite representation* of H if $V_B = V \cup E$ and $E_B = \{ve : v \in V, e \in E, v \in e\}$. In words, $B(H)$ is the graph that has a vertex for every vertex in H and a vertex for every hyperedge of H. There is an edge between a vertex $v \in V$ and a vertex $e \in E$ if v is in the hyperedge e. We say that the vertices of $B(H)$ are normal vertices if they correspond to vertices of H, and we say they are hyperedge vertices if they correspond to hyperedges of H. We picture a hypergraph H as the drawing of $B(H)$ (see Figure 1.1).

A *path* in a hypergraph H is an alternating sequence, without repetition, of vertices and hyperedges, $v_1, e_1, v_2, e_2, ...e_k, v_{k+1}$ such that for $i = 1, .., k$, e_i contains v_{i-1} and v_i.

To picture a path in a hypergraph H, we refer to the bipartite representation $B(H)$. A path in H will be a path in $B(H)$ where both end vertices correspond to vertices of H (see Figure 1.1).

Given a hypergraph H, and $v, w \in V(H)$ we say that v and w are *connected* if there exists a path between them. We say H is *connected* if any two vertices of H are connected. We say that a set $S \subseteq V(H)$ is k *connected* if for every pair of vertices from S there exist k edge disjoint paths between them. For $X \subseteq V(H)$, we say a hyperedge e *enters* X if $e \cap X \neq \emptyset$ and $e \cap (V - X) \neq \emptyset$. We define $\delta_H(X)$ to be the set of hyperedges that enter X, and $d_H(X) = |\delta_H(X)|$.

1.3 Directed hypergraphs

In this section, we present two different definitions of directed hypergraphs. We first present directed graphs.

Definition 3. *A directed graph is a pair $G = (V, E)$ such that V is a set, and E is a multiset whose members are ordered pairs of elements of V.*

We say that V is the set of vertices of D, and A is the set of arcs of D, and we denote these as $V(D)$ and $A(D)$ respectively.

More precisely, for an arc $a = (v, w)$ of D, we say that v is the tail and w is the head of a. A *dipath* in directed graph D is an alternating sequence, without repetition, of vertices and arcs, $v_1, a_1, v_2, a_2, ...a_k, v_{k+1}$ such that for $i = 1, .., k$, v_i is the tail of a_i and v_{i+1} is the head of a_i.

There are at least two natural ways to define a directed hypergraph. One way is via directed out-hypergraphs (called star hypergraphs in [1]).

Definition 4. *A* directed out-hypergraph *is a pair $D = (V, A)$ where V is a set and A is a multiset of subsets of V, and a relation which assigns to every member e of D a tail vertex tail$(e) \in e$.*

We say that the members of A are the hyperarcs of D. Let D be a directed out-hypergraph. As for directed graphs, the idea of directing a hypergraph comes from the use in paths. In this case, we have a tail vertex for every hyperarc. For a hyperarc a, we say the vertices different from tail(a) are head vertices of a. A *dipath* in D is an alternating sequence, without repetition, of vertices and hyperarcs, $v_1, a_1, v_2, a_2, ...a_k, v_{k+1}$ such that for $i = 1, .., k$, v_i is the tail of a_i and v_{i+1} is a head of a_i.

Another way for defining a directed hypergraph, used in [4], is via directed in-hypergraphs.

.

Definition 5. *A* directed in-hypergraph *is a pair $D = (V, A)$ where V is a set and A is a multiset of subsets of V, and a relation. This relation assigns to every hyperarc e of D a head vertex* head$(e) \in e$.

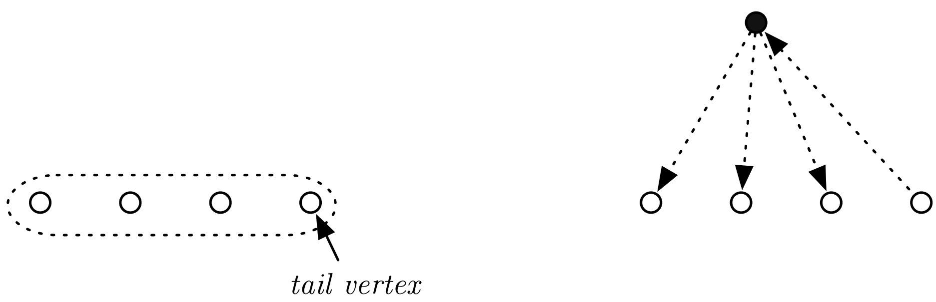

Figure 1.2: A directed out-hyperarc and its bipartite representation. (A directed in-hyperarc would simply have out-degree one at the hyperarc vertex.)

The only change is that, instead of having one tail vertex, we have one head vertex, and all non head vertices of a hyperarc are called tail vertices. The definition of a dipath is the same.

In this thesis, we are concerned with the connectivity properties that a hypergraph can maintain when orienting its hyperedges. This leads to the following definitions.

For a directed hypergraph D, if we forget about the tail (head) vertex designation of each hyperarc, we are left with a hypergraph H. We say that H is the underlying hypergraph of D. We say an *in-orientation* (*out-orientation*) of a hypergraph H is a directed in-hypergraph (out-hypergraph) D whose underlying hypergraph is H.

As for hypergraphs, we wish to picture a directed hypergraph. Let $D = (V, A)$ be a directed in-hypergraph (out-hypergraph) and H its underlying graph. We will construct a bipartite representation for D. Let $B(D)$ be the directed graph given by orienting the edges of $B(H)$ as follows. For a in D, let e be its corresponding hyperedge in H. An edge ve of $B(H)$, with $v \in e$, will be oriented as (v, e) in $B(D)$ if v is a tail of a in D. Otherwise, if v is a head of a, it will be oriented as (e, v). (We call a vertex of $B(D)$ a hyperarc vertex if it corresponds to a hyperarc of D, and we call it a normal vertex if it corresponds to a vertex of D.) See Figure 1.2.

For an orientation D of H and $X \subseteq V$, we say that a hyperarc a enters X if there is a head of a in X and a tail of a in $V - X$. We use $\delta_D^{in}(X)$ to denote the set of hyperarcs that enter X; $d_D^{in}(X) = |\delta_D^{in}(X)|$. Analogously, we say that a hyperarc arc leaves X if it enters $V - X$. We denote the set of hyperarcs that leave X as $\delta_D^{out}(X)$; and $d_D^{out}(X) = |\delta_D^{out}(X)|$. When it is not ambiguous, we omit the subscript in this notation.

We say that a directed in-hypergraph (out-hypergraph) D is *k-hyperarc connected* if it has k-hyperarc disjoint dipaths from v to w, for every pair of vertices $v, w \in V(D)$.

1.4 Menger's theorem

Consider a directed graph $D = (V, A)$. For $S, T \subseteq V$ we say that a dipath is an S,T-*dipath* if its start vertex is in S and its end vertex is in T. A set C of vertices is called S,T-*disconnecting* if C intersects each S,T-dipath (C may intersect $S \cup T$). For vertices $s, t \in V$, we say that a set of arcs C is an s,t- cut if $C = \delta^{out}(U)$ for some subset U of V with $s \in U$, t$\notin U$. For graphs, hypergraphs, and directed hypergraphs we define S,T-paths and dipaths in an analogous way. We also define an s,t-*cut* in the same way.

Menger [8], gave a min-max theorem for the maximum number of disjoint S,T paths in an undirected graph.

Theorem 1. (Menger's theorem, directed vertex-disjoint version.) *Let $D = (V, A)$ be a digraph and let $S, T \subseteq V$. Then the maximum number of vertex disjoint S,T-dipaths is equal to the minimum size of an S,T-disconnecting vertex set.*

As a corollary, we have that for vertices s, t the maximum number of arc disjoint s, t dipaths is equal to the minimum size of an $s,t-$cut.

Corollary 1. (Menger's theorem, directed arc-disjoint version.) *Let $D = (V, A)$ be a digraph and $s, t \in V$. Then the maximum number of arc disjoint s, t dipaths is equal to the minimum size of an s,t-cut.*

An edge-disjoint version of this theorem holds for undirected graphs.

Corollary 2. (Menger's theorem for hypergraphs, hyperedge-disjoint version.) *Let $H = (V, E)$ be an undirected hypergraph. Then the maximum number of hyperedge disjoint s,t-paths is equal to the minimum size of an s,t-cut.*

Proof. We construct a directed graph for which every s, t dipath corresponds to an s, t path in the hypergraph. We then apply Corollary 1.

Let $G = (V', A')$ be the directed graph with a vertex v for each $v \in V$ and two vertices x_e and y_e for every hyperedge $e \in E$. Like in the bipartite representation, we will add arcs between a hyperedge vertex and a normal vertex if the vertex is in the corresponding hyperedge. We have an arc starting in v and ending in x_e if $v \in e$. Similarly, we have an arc starting in y_e and ending in v if $v \in e$. Finally, there is an arc (x_e, y_e) for every $e \in E$ (see Figure 1.3).

Note that a dipath between two normal vertices in G corresponds to a path between the same two vertices in H and vice versa. The corollary now follows from Corollary 1.

□

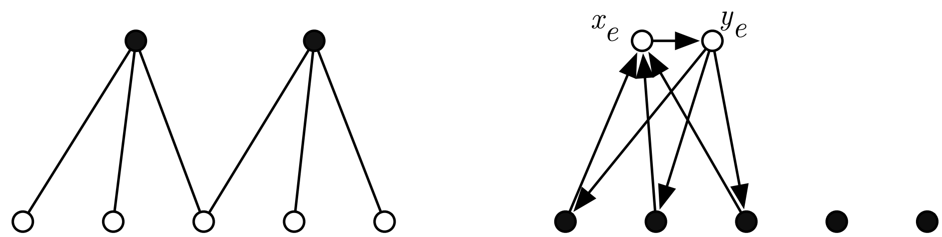

Figure 1.3: The bipartite representation $B(H)$ on the left. On the right we present D, the digraph constructed for the proof of Corollary 2. In particular, there is a one-to-one correspondence between the paths in H and the dipaths in D with end points in normal vertices.

Corollary 3. (Menger's theorem for directed in-hypergraphs (out-hypergraphs), hyperarc-disjoint version.) *Let $H = (V, A)$ be a directed in-hypergraph (out hypergraph). Then the maximum number of hyperarc disjoint s,t paths is equal to the minimum size of an s, t-cut.*

The proof of this corollary follows by applying Corollary 1 to the bipartite representation of H.

1.5 Flows and connectivity requirements

In this section we present the Edmonds-Giles theorem. We refer the reader to Schrijver's book [11] for background information. Given a ground set V, and $X, Y \subseteq V$ we say that X and Y are *intersecting* if $X \cap Y, X \backslash Y, Y \backslash X$ are all non empty. We say that a family $\mathcal{F}$ of subsets of V is *intersecting* if $X \cup Y$, $X \cap Y$ are in $\mathcal{F}$ for every pair of intersecting sets $X, Y \in \mathcal{F}$. We say that a function h on an intersecting family of subsets $\mathcal{F}$ is *intersecting supermodular* if

$$h(X) + h(Y) \leq h(X \cup Y) + h(X \cap Y)$$

for every pair of intersecting sets $X, Y \in \mathcal{F}$.

Given a hypergraph $H = (V, E)$ and an intersecting supermodular function h on $\mathcal{F} \subseteq 2^V$, we say that an orientation D of H *covers* h if $d^{in}(X) \geq h(X)$ for every $X \in \mathcal{F}$. We say that h is a *connectivity requirement* of H. Throughout this thesis all functions are integral.

For a ground set V we say that $X, Y \subseteq V$ are *crossing* if they are intersecting and $X \cup Y \neq V$. Furthermore, we say that a family of subsets $\mathcal{F}$ of V is *crossing* if $X \cup Y$, $X \cap Y$ are in $\mathcal{F}$ for every pair of crossing sets $X, Y \in \mathcal{F}$.

Let $D = (V, A)$ be a digraph and $\mathcal{C}$ be a crossing family of subsets of V. A function $f : \mathcal{C} \to \mathbb{R}$ is called *crossing submodular* if

$$h(X) + h(Y) \geq h(X \cup Y) + h(X \cap Y)$$

for every pair of crossing sets $X, Y \in \mathcal{C}$.

For a vector $x : S \to R$, and $S' \subseteq S$ we denote $\sum_{i \in S'} x_i$ by $x(S')$. Given such $D, \mathcal{C}, f$ the *submodular flow polyhedron* B consists of vectors $x \in \mathbb{R}^A$ satisfying

$$x(\delta^{in}(U)) - x(\delta^{out}(U)) \leq f(U) \text{ for each } U \in \mathcal{F}.$$

We say that a polyhedron $B \subseteq \mathbb{R}^n$ is *box-integer* if for all $c, d \in \mathbb{Z}^n$, the polytope

$$B \cap \{x \in \mathbb{R}^n | d \leq x \leq c\}$$

is integer.

Theorem 2. (Edmonds-Giles theorem.) *For a digraph $D = (V, A)$ and a crossing submodular function f on a crossing family $\mathcal{C}$ of subsets of V, the submodular flow polyhedron is box-TDI. Furthermore, if f is integer valued, the submodular flow polyhedron is box-integer.*

This theorem is used to prove many results on graph orientation problems. In this thesis, we use it to provide an algorithm for the existence of a Steiner rooted k-hyperarc connected orientation for a given indegree specification. As an example of its use, in the next section we use the Edmonds-Giles theorem to prove Nash-Williams' weak orientation theorem.

1.6 Nash-Williams' theorems

In this section we present two theorems by Nash-Williams. Both theorems are predecessors of the results presented in this thesis.

Theorem 3. (Nash-Williams' Weak Orientation Theorem) *An undirected graph G has a k-arc connected orientation if and only if G is $2k$-edge connected.*

Proof. From Menger's theorem it follows that a graph that has a k-arc connected orientation is $2k$-edge connected. It remains to prove that if G is $2k$-edge connected, then there exists an orientation that is k-arc connected. Let $D = (V, A)$ be an arbitrary orientation of G. Consider the polyhedron $\mathcal{B} \subseteq R^A$ such that $x \in \mathcal{B}$ if

$$x(\delta_D^{in}(U)) - x(\delta_D^{out}(U)) \leq d_D^{in}(U) - k, \text{ for each } U \subset V, \text{ where } U \neq \emptyset. \tag{1.1}$$

It is easy to see that $d^{in}(U) - k$ is crossing submodular. Hence, by the Edmonds-Giles Theorem (Theorem 2), $\mathcal{B}$ is box-integer. It follows that $\mathcal{B}' = \{x \in \mathcal{B} : 0 \leq x_a \leq 1, a \in A\}$ is an integral polyhedron. Note that the vector $x := \frac{1}{2} \times \mathbf{1}$ satisfies Equation 1.1 because every cut of G has size at least $2k$; hence $\mathcal{B}'$ is nonempty. Therefore, $\mathcal{B}'$ has an integer solution x. Let D' be the orientation resulting from D by reversing the orientation of the arcs a of D such that $x_a = 1$. By Menger's Theorem, D' will be a k-arc connected orientation if and only if $d^{in}_{D'}(U) \geq k$ for every $U \subset V$. From the definition of D' and the fact that x satisfies Equation 1.1, we have

$$d^{in}_{D'}(U) = d^{in}_{D}(U) - x(\delta^{in}(U)) + x(\delta^{out}(U)) \geq d^{in}_{D}(U) - (d^{in}_{D}(U) - k) \geq k.$$

□

For a graph G and an orientation D of G, let $\lambda_G(s,t)$ denote the maximum number of edge-disjoint s,t paths in G, and $\lambda_D(s,t)$ denote the maximum number of arc-disjoint s,t dipaths in D. The following is a deep result for which an analogous for hypergraphs is not known.

Theorem 4. (Nash-Williams' Strong Orientation Theorem [9]) *Any undirected graph $G = (V,E)$ has an orientation $D = (V,A)$ with*

$$\lambda_D(s,t) \geq \left\lfloor \frac{1}{2}\lambda_G(s,t) \right\rfloor, \ \forall s,t \in V. \tag{1.2}$$

1.7 Edge-disjoint spanning trees

In this section, we analyze possible extensions to hypergraphs of the following result for graphs.

Theorem 5. *(Tutte [12]; Nash-Williams [10]) An undirected graph $G = (V,E)$ has k edge-disjoint spanning trees if and only if for every partition $\mathcal{P} = \{V_1, ...V_t\}$ of V, the number of edges with end vertices in different members of $\mathcal{P}$ is at least $k(t-1)$.*

When G is a $2k$ connected graph, there will be at least $k(t-1)$ edges connecting different sets of the partition $\mathcal{P}$. That is, when G is $2k$ connected there are k edge-disjoint spanning trees. The next theorem, due to Bang-Jensen et al.[1], implies that a simple analogue of this fact does not hold for hypergraphs.

Theorem 6. *[1] For every positive integer k there exists an undirected hypergraph H that is k-hyperedge connected, and does not contain two hyperedge-disjoint spanning connected subhypergraphs.*

Nevertheless, Frank et al. [5] proved an analogue to Tutte's theorem.

We say that a hypergraph is *k-partition connected* if the number of hyperedges with endpoints in at least two members of $\mathcal{P}$ is at least $k(t-1)$, where t is the number of members of $\mathcal{P}$.

Theorem 7. *[5] An undirected hypergraph is k-partition connected iff it can be decomposed into k hyperedge-disjoint 1-partition connected subhypergraphs.*

Note that a connected hypergraph may have just one hyperedge, but a 1-partition connected hypergraph has at least $|V|-1$ hyperedges (consider the partition into singletons, see Figure 1.4).

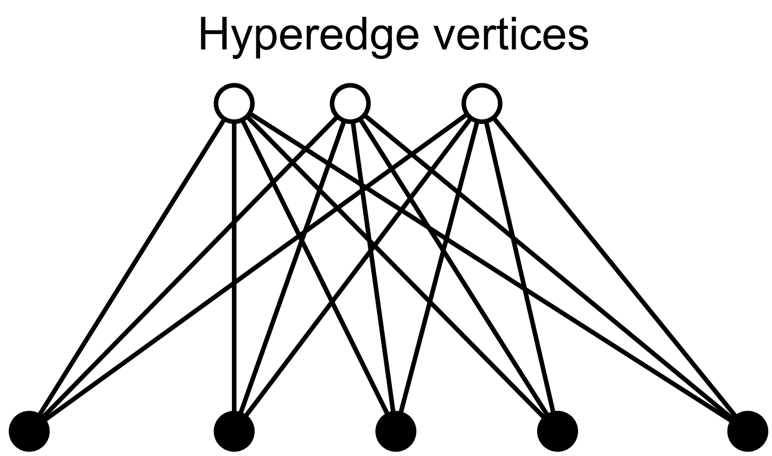

Figure 1.4: The bipartite representation of a 3-hyperedge connected hypergraph. Note that it has only 3 hyperedges. By taking the partition of the vertices consisting of the singletons, it follows that this hypergraph is not 1-partition connected.

1.8 Steiner graphs and Steiner hypergraphs

Throughout this thesis, we work with Steiner hypergraphs. For a hypergraph $H=(V,E)$ and a designated set of vertices $S \subseteq V$, we say that the *Steiner vertices* are the vertices in $V-S$. We say that a vertex in S is a *terminal vertex*, and sometimes we refer to it as a terminal.

It is natural to try to extend Theorem 5 or theorems like the Nash-Williams orientation theorems to the setting of Steiner graphs. Unfortunately, most of these problems on Steiner graphs are NP-hard and so there exist no good characterizations or min-max theorems, assuming NP$\neq$ co-NP.

1.9 Reversing paths

A key operation that will be used in the thesis is what we call *reversing* a path.

Consider a hypergraph $H = (V, E)$ and an out-orientation $D = (V, A)$ of H. Let $\overrightarrow{P} = (v_1, a_1, ..., a_k, v_{k+1})$ be a dipath in D, where e_i is the corresponding hyperedge in H of a_i. We say that D' is the orientation obtained from D by *reversing* P if D' is an out-orientation of H, where the tails are assigned in the same way as in D except for $e_1, ...e_k$, and for $i = 1, ..., k$, we assign e_i to have tail v_{i+1}; thus v_i is a head of e_i. Note that $\overleftarrow{P} = (v_{k+1}, e_k, v_k, e_{k-1}, ..., e_1, v_1)$ is a dipath of D'.

1.10 Some results on in-orientations

The *hypergraph orientation problem* is to find an orientation of a hypergraph that covers a given connectivity requirement function. We say that an orientation of a hypergraph is *out-rooted k-hyperarc connected* if there exists a vertex r such that there are k-hyperarc disjoint dipaths from r to every other vertex. We say that it is *in-rooted k hyperarc connected* if there exists a vertex r such that there are k hyperarc disjoint dipaths from every other vertex to r.

Most results that can be proved for in-orientations can be translated to equivalent results for out-orientations, and vice versa. We do not mean that the same result will hold for both, but that any in-orientation can be reversed to give an out-orientation, where the ‘opposite” result will hold. For example, if a hypergraph admits an out-orientation that is out-rooted k-hyperarc connected (k paths from the root to each vertex), then it admits an in-orientation that is in-rooted k-hyperarc connected (k paths from each vertex to the root).

In [3], Frank studies the graph orientation problem. He gives a characterization for intersecting supermodular connectivity requirements. Two decades later, Frank et al. [4] studied the hypergraph orientation problem, this time using crossing supermodular connectivity requirements, and they proved the following theorem.

Theorem 8. *Let $H = (V, E)$ be an undirected hypergraph, and let h be a non-negative crossing supermodular set function. There is an in-orientation of H covering h if and only* *if*

$$\sum_{X\in\mathcal{P}} h(X) \le |\{e \in E : e \in \delta(X), X \in \mathcal{P}\}|,$$
$$\sum_{X\in\mathcal{P}} h(V-X) \le \sum_{e\in E} \max\{0, (|\{X \in \mathcal{P} : e \in \delta(X)\}| - 1)\}$$

for every partition $\mathcal{P}$ of V.

Theorem 8 can be used to prove that, for any $2k$-hyperedge connected hypergraph H and a root vertex $r \in V(H)$, H has an in-rooted k-hyperarc connected in-orientation. This is done by defining appropriate connectivity requirements.

For non-negative integers $k \ge l$, a directed hypergraph is called (k,l)-*edge-connected* if there is a node $s \in V$ such that there are k edge-disjoint dipaths from s to every other node, and there are l edge-disjoint dipaths to s from every other node. For non-negative integers $k \ge l$, a hypergraph H is called (k,l)-*partition-connected* if the number of edges that have end vertices in different members of $\mathcal{P}$ is at least $k(t-1)+l$ for every partition $\mathcal{P}$ with t members. More precisely, if $|\{e \in E : \exists X \in \mathcal{P}, e \in \delta(X)\}| \ge k(t-1)+l$. When $l = 0$ we simply call it k-partition connected.

Corollary 4. *[4] An undirected hypergraph has a (k, l)-edge-connected in-orientation if and only if it is (k,l)-partition-connected.*

Chapter 2

Orienting Hypergraphs

2.1 Introduction

In this chapter, we present a result by A. Frank [3] in the setting of hypergraphs. Frank's theorem characterizes graphs that can be oriented to satisfy certain connectivity requirements.

The main result of this chapter, Theorem 10, is similar to Theorem 8 with three main differences. The first one is that Theorem 10 allows for negative values of the requirement function. This directly gives a corollary on orienting mixed hypergraphs (a hypergraph with directed hyperedges and undirected hyperedges). Another difference is that in Theorem 10 the requirement function h does not have to be defined on all subsets of V. Finally, the functions for Theorem 10 are a bit more general than supermodular functions.

Recall from Chapter 1, that $d_G(X, Y)$ denotes the number of edges between $X \backslash Y$ and $Y \backslash X$.

Definition 6. *Given a graph $G = (V, E)$ and an intersecting family $\mathcal{F}$ of subsets of V, we say that $f : \mathcal{F} \to \mathbb{Z}$ is* relaxed supermodular *if for every two intersecting sets X and Y from $\mathcal{F}$,*

$$f(X) + f(Y) \leq f(X \cup Y) + f(X \cap Y) + d_G(X, Y).$$

We say that $\mathcal{P} = \{V_1, ..., V_t\}$ is a *subpartition* of V if $V_i \subseteq V$ for all $i = 1, ..., t$, and furthermore, if $i \neq j$, then $V_i \cap V_j = \emptyset$. For a family $\mathcal{F}$ of subsets of a ground set V, we say that a subpartition $\mathcal{P} = \{V_1, ..., V_t\}$ of V is a *subpartition from* $\mathcal{F}$ if every member of $\mathcal{P}$ is in $\mathcal{F}$. For a subpartition $\mathcal{P}$ of V, we use $e(G, \mathcal{P})$ to denote the number of edges that either have end vertices in different members of $\mathcal{P}$ or have one vertex in $\cup\{V_i : V_i \in \mathcal{P}\}$, that is $e(G, \mathcal{P}) = |\{e \in E : \exists V_i \in \mathcal{P}, e \in \delta(V_i)\}|$.

Theorem 9. *([3], Theorem 2] For an undirected graph $G = (V, E)$, an intersecting family $\mathcal{F}$ of subsets of V such that $\emptyset \notin \mathcal{F}, V \in \mathcal{F}$, and a relaxed supermodular function $f : \mathcal{F} \to \mathbb{Z}$ such that $f(V) = 0$, G has an orientation that covers f if and only if*

$$e(G, \mathcal{P}) \geq \sum_{V_i \in \mathcal{P}} f(V_i) \tag{2.1}$$

for every subpartition $\mathcal{P} = \{V_1, ..., V_t\}$ from $\mathcal{F}$.

Note that, Frank uses the term convex functions to denote relaxed supermodular functions. Instead of saying that $\mathcal{F}$ is an intersecting family of subsets of V, he explicitly adds the definition to the theorem. Besides the notational differences, this is the same theorem presented by Frank in [3].

In this chapter, we extend Theorem 9 to out-orientations of hypergraphs. More precisely, for a hypergraph H and a requirement function f, we state necessary and sufficiency conditions for the existence of an out-orientation of H that covers f. In this chapter, we refer to directed out-hypergraphs simply as directed hypergraphs, and out-orientations simply as orientations.

In Section 2 we introduce notation specific to this chapter. In Section 3 we state the main theorem of this chapter and discuss some applications. Section 4 presents an overview of the proof together with some results about tight sets. Section 5 presents the proof of the main theorem.

2.2 Preliminaries

In this section, we extend some notation, commonly used in graphs, to hypergraphs.

2.2.1 Notation

We recall some terms and notation from Chapter 1. Given a hypergraph $H = (V, E)$ and $X \subseteq V$, we say a hyperedge e *enters* X, if $e \cap X \neq \emptyset$ and, at the same time, $e \cap (V - X) \neq \emptyset$. We define $\delta_H(X)$ to be the set of hyperedges that enter X, and $d_H(X) = |\delta_H(X)|$. For an orientation D of H and $X \subseteq V$, we say that a hyperarc a *enters* X if there is a head of of a in X and a tail of a in $V - X$. We use $\delta_D^{in}(X)$ to denote the set of hyperarcs that enter X; $d_D^{in}(X) = |\delta_D^{in}(X)|$. Analogously, we say that a hyperarc arc *leaves* X if it enters $V - X$. We denote the set of hyperarcs that leave X as $\delta_D^{out}(X)$; and $d_D^{out}(X) = |\delta_D^{out}(X)|$.

As we have already mentioned, in this chapter we work only with directed out-hypergraphs, and we refer to them simply as directed hypergraphs. Recall that an out-orientation D of

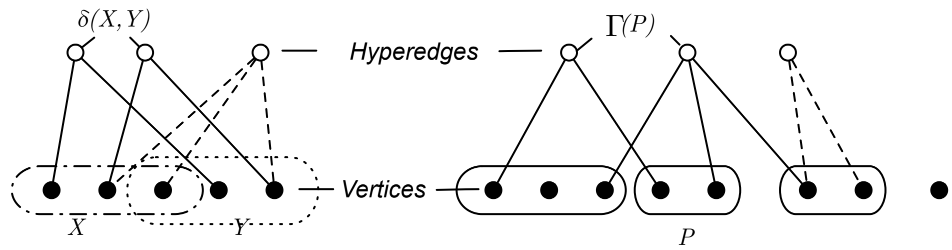

Figure 2.1: On the left, the bipartite representation of a hypergraph, subsets of vertices X, Y, and the hyperedges (hyperedge vertices) in $\delta(X, Y)$. On the right, the bipartite representation of a hypergraph, a subpartition $\mathcal{P}$, and the hyperedges (hyperedge vertices) in $\Gamma(\mathcal{P})$

a hypergraph H, assigns to every hyperedge of H a tail vertex. For a hyperarc $a \in A(D)$ we use $\mathrm{tail}(a)$ to denote the tail vertex of a.

When extending Theorem 9 to hypergraphs, some difficulties arise. Namely, $d(X, Y)$ and $e(G, \mathcal{P})$ are not well defined for hypergraphs. We extend the definitions of $d(X, Y)$ and $e(G, \mathcal{P})$ to the setting of hypergraphs in such a way that the meaning of these symbols for graphs remains the same.

Given a hypergraph $H = (V, E)$ and $X, Y \subseteq V$, we let $\delta_H(X, Y)$ denote the set of hyperedges that enter $Y \backslash X$ and $X \backslash Y$ but do not enter $X \cap Y$; $d_H(X, Y) = |\delta_H(X, Y)|$. Note that $d(X, Y)$ is defined only for hypergraphs, and not for directed hypergraphs. If D is an orientation of H, we say that $\delta_D(X, Y) = \delta_H(X, Y)$; $d_D(X, Y) = d_H(X, Y)$. See Figure 2.1.

Lemma 1. *Let $H = (V, E)$ be a hypergrah, and D be an orientation of H. Then, for $X, Y \subseteq V$, we have*

$$d^{in}_D(X) + d^{in}_D(Y) = d^{in}_D(X \cup Y) + d^{in}_D(X \cap Y) + d_D(X, Y). \tag{2.2}$$

Proof. Each term in the equation is the cardinality of a set of hyperarcs. Hence it suffices to show that every hyperarc of D will appear the same number of times in the sets corresponding to the values of LHS as of the sets corresponding to the values of the RHS.

For the inequality $d^{in}_D(X) + d^{in}_D(Y) \le d^{in}_D(X \cup Y) + d^{in}_D(X \cap Y) + d_D(X, Y)$: Let a be a hyperarc of D. If a enters X but not Y, then either a enters $X \cup Y$, a enters $X \cap Y$ or a is in $\delta_D(X, Y)$. So a appears on the RHS at least the same number of times as it appears in the LHS. The case where a enters Y, but not X, is analogous. If a enters both X and Y, its tail is not in $X \cup Y$ so it must enter $X \cup Y$, and it either enters $X \cap Y$, or it is in $\delta_D(X, Y)$.

For $d_D^{in}(X) + d_D^{in}(Y) \geq d_D^{in}(X \cup Y) + d_D^{in}(X \cap Y) + d_D(X, Y)$: First consider a hyperarc a that enters $X \cup Y$. Without loss of generality, we can assume that a enters X. Note that a cannot enter $X \cap Y$ and be in $\delta_D(X, Y)$ at the same time. If a enters $X \cap Y$ then a enters Y. If a is in $\delta_D(X, Y)$, then a must enter Y. Now consider a hyperarc a that does not enter $X \cup Y$. If the right-hand-side is zero then we are done. Therefore we may assume that either a enters $X \cap Y$ or a is in $\delta_D(X, Y)$, but not both. If a enters $X \cap Y$, then either a enters X, or a enters Y. If a is in $\delta_D(X, Y)$, then it either enters X or it enters Y. □

Definition 7. *For an undirected hypergraph $H = (V, E)$, and a subpartition $\mathcal{P} = \{V_1, ...V_t\}$ of V, we say that a hyperedge e is* mishandled *by $\mathcal{P}$ if e is contained in $\cup_{i=1}^{t} V_i$ and e enters at least one of the members of $\mathcal{P}$.*

Let $\Gamma(\mathcal{P})$ denote the set of hyperedges that are mishandled by P, and let $\gamma(\mathcal{P}) = |\Gamma(\mathcal{P})|$. We define $e_H(\mathcal{P}) = (\sum_{V_i \in \mathcal{P}} d(V_i)) - \gamma(\mathcal{P})$. See Figure 2.1. When it is not ambiguous, we simply use $e(\mathcal{P})$.

Note that a graph $G = (V, E)$ is a hypergraph with hyperedges of size two (edges), $e(G, \mathcal{P}) = |\{e \in E : \exists V_i \in \mathcal{P}, e \in \delta(V_i)| = (\sum_{V_i \in \mathcal{P}} d(V_i)) - \gamma(\mathcal{P}) = e(\mathcal{P})$. The definition of $e(\mathcal{P})$ allows us to relate the connectivity properties of a hypergraph with the connectivity properties of an orientation of the same hypergraph.

The following proposition differentiates an out-orientation of a hypergraph from other ways of defining an orientation; if D is an orientation such that the number of tail and head nodes is arbitrary, then an equality comparing d_D^{in} and d_H would be more difficult, if even possible, to obtain.

Proposition 1. *Let $D = (V, A)$ be an orientation of a hypergraph H. Let C be a subset of V such that no hyperarcs from D leave C, and let $\mathcal{P} = \{V_1, ..., V_t\}$ be a partition of C. Then*

$$e(\mathcal{P}) = \sum_{V_i \in \mathcal{P}} d_D^{in}(V_i) \tag{2.3}$$

Proof. For a hyperedge e and its corresponding hyperarc $\overrightarrow{e}$, e contributes to $e(\mathcal{P})$ as much as $\overrightarrow{e}$ contributes to $\sum_{V_i \in \mathcal{P}} d_D^{in}(V_i)$. If tail($\overrightarrow{e}$) is in C, e is either mishandled by $\mathcal{P}$ or it is not mishandled by $\mathcal{P}$. In the latter case neither e nor $\overrightarrow{e}$ contribute to the equation because e is contained in a member V_i of $\mathcal{P}$. So assume that e is mishandled by the partition, then e contributes one more to $\sum_{V_i \in \mathcal{P}} d_H(V_i)$ than $\overrightarrow{e}$ to $\sum d_D^{in}(V_i)$. On the other hand, e also contributes one to $\gamma(\mathcal{P})$, and this is subtracted from $e(\mathcal{P})$. If tail($\overrightarrow{e}$) is not in C, then e contributes to $\sum_{V_i \in \mathcal{P}} d_H(V_i)$ the same number as the contribution of $\overrightarrow{e}$ to $\sum d_D^{in}(V_i)$. □

For a hypergraph, we define a relaxed supermodular function in the same way as for graphs.

Definition 8. *Given an undirected hypergraph $H = (V, E)$ and an intersecting family $\mathcal{F}$ of subsets of V, we say that $f : \mathcal{F} \to \mathbb{Z}$ is relaxed supermodular if*

$$f(X) + f(Y) \leq f(X \cup Y) + f(X \cap Y) + d_H(X, Y), \tag{2.4}$$

for all intersecting pairs $X, Y \in \mathcal{F}$.

2.3 The main theorem and its applications

The main result of this chapter is the following characterization of hypergraphs that admit orientations covering relaxed supermodular connectivity requirements.

Theorem 10. *Let $H = (V, E)$ be an undirected hypergraph and $\mathcal{F}$ an intersecting family from V such that $\emptyset \notin \mathcal{F}, V \in \mathcal{F}$. Let f be a relaxed supermodular function on $\mathcal{F}$ such that $f(V) = 0$. Then H has an orientation that covers f if and only if, for every subpartition $\mathcal{P} = \{V_1, ...V_t\}$ from $\mathcal{F}$,*

$$e(\mathcal{P}) \geq \sum_{V_i \in \mathcal{P}} f(V_i). \tag{2.5}$$

We present two applications of Theorem 10. For the first application, we first define a mixed hypergraph.

Definition 9. *We say that $H = (V, E, \overrightarrow{A})$ is a mixed hypergraph if E is a multiset of hyperedges of V and $\overrightarrow{A}$ is a multiset of hyperarcs of V.*

For a mixed hypergraph $H = (V, E, \overrightarrow{A})$, we have the associated hypergraphs $H_E = (V, E)$ and $H_A = (V, A)$, where A is the set of hyperedges induced by the hyperarcs in $\overrightarrow{A}$. We also have an associated directed hypergraph $D_{\overrightarrow{A}} = (V, \overrightarrow{A})$. Let $e_E(\mathcal{P})$ denote $e_{H_E}(\mathcal{P})$, $d_{\overrightarrow{A}}(X, Y)$ be $d_{H_A}(X, Y)$ and $d^{in}_{\overrightarrow{A}}$ be $d^{in}_{D_A}(X)$. We have the following extension of Theorem 10 for mixed hypergraphs.

Corollary 5. *Suppose $H = (V, E, \overrightarrow{\mathcal{A}})$ is a mixed hypergraph. Let $\mathcal{F}$ be an intersecting family of subsets from V such that $\emptyset \notin \mathcal{F}, V \in \mathcal{F}$. Let f be a relaxed supermodular function on $\mathcal{F}$ with $f(V) = 0$. Then H has an orientation that covers f if and only if, for every subpartition $\mathcal{P} = \{V_1, ...V_t\}$ from $\mathcal{F}$,*

$$e_E(\mathcal{P}) \geq \sum_{V_i \in \mathcal{P}} f'(V_i) \tag{2.6}$$

where $f' : \mathcal{F} \to \mathbb{Z}$ is given by $f - d^{in}_{\overrightarrow{\mathcal{A}}}$.

Proof. Notice that f' is relaxed supermodular on $\mathcal{F}$: for intersecting sets X,Y from $\mathcal{F}$, recall that Lemma 1 states that $d^{in}_{\vec{\mathcal{A}}}(X) + d^{in}_{\vec{\mathcal{A}}}(Y) = d^{in}_{\vec{\mathcal{A}}}(X \cup Y) + d^{in}_{\vec{\mathcal{A}}}(X \cap Y) + d_{\vec{\mathcal{A}}}(X, Y)$. By changing the signs and deleting $d_{\vec{\mathcal{A}}}(X, Y)$ from the inequality, we obtain that $-d^{in}_{\vec{\mathcal{A}}}(X) - d^{in}_{\vec{\mathcal{A}}}(Y) \leq -d^{in}(X \cup Y)_{\vec{\mathcal{A}}} - d^{in}_{\vec{\mathcal{A}}}(X \cap Y)$. This, together with the relaxed supermodularity of f implies that f' is relaxed supermodular, and the corollary follows from Theorem 10. $\square$

Essentially, this corollary allows us to cope with mixed hypergraphs, where the hypearcs are out-oriented. If the hyperarcs of the mixed graph are not out-oriented, then we cannot guarantee that f' is relaxed supermodular.

Another application of Theorem 10 pertains to the rooted hypergraph orientation problem. In this problem, we are given a hypergraph, and we wish to find an orientation such that the connectivity from a given root vertex to the rest of the vertices is maximized.

Corollary 6. *Let $H = (V, E)$ be an undirected hypergraph. Fix an arbitrary vertex r to be the root vertex. If H is 2k-hyperedge connected, then there exists an orientation of H such that there are k hyperarc disjoint dipaths from r to every other vertex.*

Proof. Let $\mathcal{F}= 2^V - \emptyset$. Define a function $f : \mathcal{F} \to \mathbb{Z}$ by

$$f(X) = \begin{cases} k & \text{if } r \notin X \\ 0 & \text{otherwise.} \end{cases} \tag{2.7}$$

By Menger's Theorem, an orientation that covers f has k hyperarc disjoint dipaths from r to every other vertex. Notice that f is relaxed supermodular. By Theorem 10, there exists an orientation covering f if and only if $e(\mathcal{P}) \geq \sum_{V_i \in \mathcal{P}} f(V_i)$ for every subpartition $\mathcal{P} = \{V_1, ...V_t\}$ of V. We need to show that the inequality holds for every subpartition.

Let $\mathcal{P} = \{V_1, ..., V_t\}$ be a subpartition such that $r \notin \cup_{i=1}^{t}(V_i)$, hence $\sum_{V_i \in \mathcal{P}} f(V_i) = kt$. On the other hand, because H is $2k$-hyperedge connected, $d(V_i) \geq 2k$. Note that for a hyperedge to be in $\Gamma(\mathcal{P})$, it has to enter at least two members of $\mathcal{P}$, hence $\gamma(\mathcal{P}) \leq \sum_{V_i \in \mathcal{P}} d(V_i)/2$. Taking all this into account, $e(\mathcal{P}) \geq \sum_{V_i \in \mathcal{P}} d(V_i)/2 \geq 2kt/2 = kt$, thus $e(\mathcal{P}) \geq \sum_{V_i \in \mathcal{P}} f(V_i)$.

If $r \in \cup_{i=1}^{t}(V_i)$, then $\sum_{V_i \in \mathcal{P}} f(V_i) = k(t-1)$, nevertheless, the same analysis shows that $e(\mathcal{P}) \geq k(t-1)$.

$\square$

2.4 Overview of the proof

In this section, we give a high-level description of the proof of Theorem 10. We also discuss some small, but important propositions that are used throughout the proof.

2.4.1 Overview

For the proof of Theorem 10 we follow Frank's proof of the corresponding theorem for graphs [3]. The proof consists of an induction over $val(f) = \sum\{f(X) : X \in \mathcal{F}, f(X) \geq 0\}$. We construct a relaxed supermodular function $f' : \mathcal{F} \to \mathbb{Z}$ such that $f' \leq f$ and $val(f') < val(f)$. By the induction hypothesis, there exists an orientation D of H that covers f'. We then modify this orientation so that it covers f. For this, we construct a directed hypergraph $D(f')$ by adding some arcs to D. In this new directed hypergraph, we find a set of hyperarcs whose reorientation gives us an orientation that covers f.

2.4.2 Tight sets

First, we study some properties of orientations that cover f. Let H, $\mathcal{F}$, f be as in Theorem 10. Let D be an orientation of H that covers f. A set of terminals X is said to be *tight* with respect to f if $d_D^{in}(X) = f(X)$. Note that a set may be tight with respect to an orientation and not with respect to undirected hypergraphs.

In the proof of Theorem 10, the following two propositions might be used without mention.

Proposition 2. *The union and the intersection of every two intersecting tight sets is tight.*

Proof. Let X and Y be two intersecting tight sets. The first and last term of the inequality sequence below are equal:

$$\begin{aligned} d_D^{in}(X) + d_D^{in}(Y) &= f(X) + f(Y) \\ &\leq f(X \cup Y) + f(X \cap Y) + d_D(X, Y) \\ &\leq d_D^{in}(X \cup Y) + d_D^{in}(X \cap Y) + d_D(X, Y) \\ &= d_D^{in}(X) + d_D^{in}(Y) \end{aligned}$$

This means that

$$f(X \cup Y) + f(X \cap Y) + d_D(X, Y) = d_D^{in}(X \cup Y) + d_D^{in}(X \cap Y) + d_D(X, Y),$$

hence $f(X \cup Y) = d_D^{in}(X \cup Y)$ and $f(X \cap Y) = d_D^{in}(X \cap Y)$. That is, $X \cup Y$ and $X \cap Y$ are tight. $\square$

Proposition 3. *The intersection of a family of tight sets is a tight set, provided it is nonempty. If a family of tight sets forms a connected hypergraph, then their union is tight.*

Proof. Both statements follow by a repeated application of Proposition 2. $\square$

2.5 The proof

We present the proof of Theorem 10.

2.5.1 The induction

We use induction on $val(f) = \sum\{f(X) : X \in \mathcal{F}, f(X) \geq 0\}$. It is straightforward that every orientation is good when $val(f) = 0$. More precisely, if $val(f) = 0$, any orientation covers f.

Assume that $val(f) > 0$. Our goal is to construct a function $f' : \mathcal{F} \to \mathbb{Z}$ such that $f' \leq f$ and $val(f') < val(f)$. Let X be a set of terminals X such that $f(X) > 0$. Then from Equation 2.5 it follows that there exists a hyperedge Z that enters X. Let b and a be in Z such that b is in X and a is not. We say that a subset Y of V is a $b\overline{a}$ set if $b \in Y$ and $a \notin Y$. Define a function f' on $\mathcal{F}$ as follows:

$$f'(Y) = \begin{cases} f(Y) - 1 & \text{if } Y \text{ is } b\overline{a} \text{ set} \\ f(Y) & \text{otherwise.} \end{cases} \tag{2.8}$$

2.5.2 About f'

Note that f' is relaxed supermodular with respect to $\mathcal{F}$. Furthermore, $val(f') < val(f)$. Hence we can use the induction hypothesis. Let D be an orientation of H that covers f'. Assume that D does not cover f. Then there exists a $b\overline{a}$ set X with $d_D^{in}(X) = f'(X)$.

For the rest of this chapter, unless specified, we refer to a tight set, as tight with respect to f' and the above D. The following is a neat property of tight sets.

Proposition 4. *Any tight $a\overline{b}$ set A and any tight $b\overline{a}$ set B are disjoint.*

Proof. For a contradiction, suppose $A \cap B \neq \emptyset$. Because $b \notin A$ and $a \notin B$, A and B must be intersecting. By the inequalities below, we derive that $d^{in}(A) + d^{in}(B) < d^{in}(A \cup B) + d^{in}(A \cap B) + d(X, Y)$, contradicting Lemma 1.

The series of inequalities below is derived from the fact that D covers f', the fact that f is relaxed supermodular, and that f' and f have the same values on $A \cup B$ and $A \cap B$

(since neither $A \cup B$ nor $A \cap B$ are $b\overline{a}$ sets).

$$\begin{aligned} d_D^{in}(A) + d_D^{in}(B) &= f'(A) + f'(B) \\ &= f(A) + f(B) - 1 \\ &< f(A \cup B) + f(A \cap B) + d_H(X,Y) \\ &= f'(A \cup B) + f'(A \cap B) + d_H(X,Y) \\ &\leq d_D^{in}(A \cup B) + d_D^{in}(A \cap B) + d_H(X,Y). \end{aligned}$$

This shows that $d_D^{in}(A) + d_D^{in}(B) < d_D^{in}(A \cup B) + d_D^{in}(A \cap B) + d_H(X,Y)$, which is a contradiction to Lemma 1.

$\square$

2.5.3 How do we modify the orientation?

Because D covers f', it is natural to try to add an incoming hyperarc to every $b\overline{a}$ tight set. If the smallest tight set containing b also contains a, then there is no b$\overline{a}$ tight set. Hence D covers f. Therefore, it is natural to try an induction over the size of the smallest tight set containing b.

Let $T(x)$ denote the intersection of all tight sets containing a vertex x. This is well defined since V is tight. Proposition 3 tells us that the intersection of tight sets is tight, hence $T(x)$ is tight. Furthermore, $T(x)$ is the smallest tight set containing x.

As we have already mentioned, we can assume that $T(b)$ does not contain a, otherwise D covers f. From Equation 2.5 it follows that $f'(T(b)) = f(T(b)) - 1 < d(T(b))$. Hence there exists a hyperarc $\overrightarrow{Z}$ leaving $T(b)$. Let y be the tail of $\overrightarrow{Z}$ (then $y \in T(b)$) and let z be a head of $\overrightarrow{Z}$ that is not in $T(b)$. Let D' be the orientation resulting from D, by reorienting the hyperarc $\overrightarrow{Z}$ to have tail z, and denote the resulting hyperarc as $\overrightarrow{Z'}$. Hence, $\overrightarrow{Z'}$ enters $T(b)$. The smallest tight set containing b in D' would no longer be $T(b)$. If D' covers f', we could try an induction over the size of $T(b)$. Unfortunately, there is no guarantee that D' covers f'. When we reorient $\overrightarrow{Z}$ the sets that contain z may lose an incoming hyperarc. However, the only sets where the connectivity requirements might fail are the tight sets that contain z.

We could repeat the process, reorienting an outgoing arc of $T(z)$, to be incoming. The question remains, when will we stop? If at some point, we reorient a hyperarc to have tail in a vertex c, for which $T(c)$ contains both a and b, it seems that we could stop the process. It turns out that such a vertex c does exist. Furthermore, if we get such a vertex c, where the least number of reorientations is needed, then the new orientation will cover f.

The following directed hypergraph captures the process just mentioned. Let $D(f')$ be the directed hypergraph obtained from D by adding the following arcs, which we call red arcs (these are simply hyperarcs consisting of two vertices). Add a red arc from x to every other vertex of $T(x)$, for all $x \in V$. Note that there is no red arc leaving a tight set, since the opposite would contradict the minimality of $T(x)$ for some $x \in V$.

The process for finding a new orientation that covers f can be stated in a simple way. We need to find a dipath in the extended hypergraph $D(f')$, such that when reversing the dipath and deleting the red arcs, the orientation of H that remains covers f.

Let C denote the set of vertices which can be reached from b by a dipath in $D(f')$.

Observation 1. *There is no red arc, nor hyperarc, leaving C. Furthermore, C is the union of tight sets $T(x)$, $x \in C$.*

If there was a red arc leaving C, then there would be a vertex not in C, that can be reached from b, which would be a contradiction. This also implies that $\cup_{x \in C} T(x) \subseteq C$. The other containment is trivial, since $x \in T(x)$ for every vertex.

Proposition 5. *There exists a vertex x in C for which $a, b \in T(x)$*

Proof. For a contradiction, suppose there is no such vertex x. Consider the hypergraph formed by the vertices of C and the hyperedge given by the sets $T(x)$, $x \in C$. By Proposition 3, the components of this hypergraph partition C into tight sets $V_1, ...V_t$, where $V_i \in \mathcal{F}$, for all $i = 1, .., t$.

Assume that a,$b \in V_i$ for some $i \in \{1, ..., t\}$. Then there exists a sequence of hyperedges $X_1, ...X_s$, such that $a \in X_1$, $b \in X_s$ and $X_i \cap X_{i+1} \neq \emptyset$. Let this sequence be of minimum length. Because we are assuming that there is no vertex $x \in C$ such that $a, b \in T(x)$, then $s > 1$. Hence $A = \cup_{i=1}^{s-1} X_i$ (an $a\bar{b}$ set) and $B = X_s$ (a $b\bar{a}$ set) are not disjoint. This contradicts Proposition 4. Therefore, a and b must be in different components.

Let $\mathcal{P} = \{V_1, ..., V_t\}$. Without loss of generality, assume that $b \in V_1$ and a is not. Then $f'(V_1) = f(V_1) - 1$, and because every member of $\mathcal{P}$ is tight it follows that $f(V_i) = f'(V_i) = d_D^{in}(V_i)$ for every $V_i \in \mathcal{P} - V_1$. Remember that there are no hyperarcs leaving C. By Proposition 1, $e(\mathcal{P}) = \sum d_D^{in}(V_i) = \sum f'(V_i) = \sum f(V_i) - 1$, which contradicts Theorem 10's hypothesis, that is, Equation 2.5. □

Let P be a dipath in $D(f')$, from b to a vertex c, such that $a, b \in T(c)$ and the length of P is as small as possible. We shall often use the property that a tight set containing c also contains both a and b.

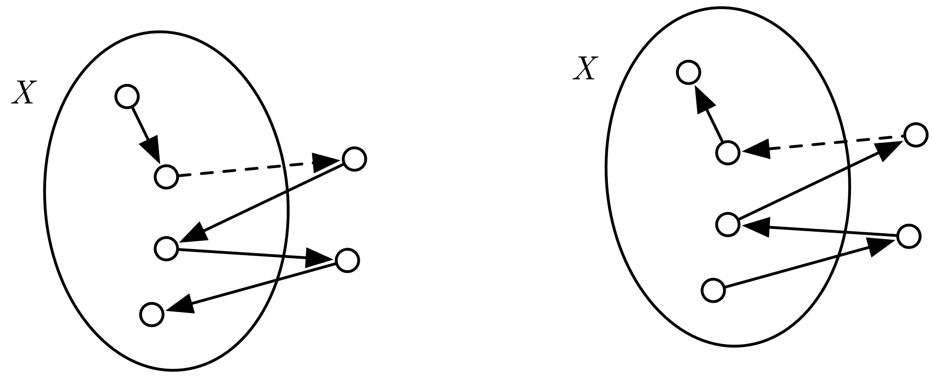

Figure 2.2: On the left, a dipath P with a red arc, represented by a dotted arrow. On the right, dipath P reversed. Note that $d^{in}(X)$ decreases by one. In the general case, $d^{in}(X)$ decreases by at most $d^{out}_{P\,\mathrm{red}}(X)$ (without taking $\epsilon(X)$ into account).

2.5.4 Reversing P fixes the orientation

Lemma 2. *Let D' be the orientation of H obtained by reversing P in $D(f')$ and deleting the red arcs of $D(f')$. Then D' covers f.*

Let ϵ be a function defined on $\mathcal{F}$ as follows:

$$\epsilon(X) = \begin{cases} -1 & \text{if X is a } c\bar{b}\text{-set} \\ +1 & \text{if X is a } b\bar{c}\text{-set} \\ 0 & \text{otherwise} \end{cases}$$

Observation 2. *Let $d^{out}_{P\,red}(X)$ denote the number of red arcs from P leaving X. Then, $d^{in}_{D'}(X) \geq d^{in}_{D}(X) + \epsilon(X) - d^{out}_{P\,red}(X)$.*

If there were no red arcs in P leaving X, then the inequality would be straight forward. In this case, $d^{in}_{D'}(X)$ would be one more than $d^{in}_{D}(X)$ in $b\bar{c}$ sets, one less in $c\bar{b}$ sets, and otherwise it would be the same. If there are red arcs in P leaving X, when we reverse the hyperarcs from P, then X might have fewer incoming arcs in D' than in D. However, the number of incoming hyperarcs to X decreases by at most $d^{out}_{P\,\mathrm{red}}(X)$ (see Figure 2.2).

To prove Lemma 2 we only need to prove that

$$d^{in}_{D}(X) + \epsilon(X) - d^{out}_{P\,\mathrm{red}}(X) \geq f(X). \tag{2.9}$$

2.5.5 The slack is enough

It is easy to check that $\epsilon(X)+\epsilon(Y)=\epsilon(X\cap Y)+\epsilon(X\cup Y)$. This with Lemma 1 and the fact that f is relaxed supermodular show that $\alpha(X)=d_D^{in}(X)-f(X)+\epsilon(X)$ is intersecting submodular with respect to $\mathcal{F}$, that is $\alpha(X)+\alpha(Y)\geq\alpha(X\cap Y)+\alpha(X\cup Y)$. Think of $\alpha(X)$ as the amount of slack that a set has in D, that is the number of incoming hyperarcs a set can afford to lose, before the orientation is no longer good in this set. We now prove that the slack of a set is always greater or equal than the number of incoming hyperarcs that it might lose when reversing P.

Proposition 6. $d^{out}_{P\ red}(X)\leq\alpha(X)$

Note that this proposition proves Lemma 2 and with it Theorem 10.

Proof. We use induction on $d^{out}_{P\,\mathrm{red}}(X)$. For this, we first prove that $\alpha(X)\geq 0$ for all $X\subset V$. This is equivalent to proving that $d_D^{in}(X)\geq f(X)-\epsilon(X)$.

If X is a $c\bar{b}$ set, because $T(c)$ is a subset of every tight set containing c, X is not tight. Thus, $d^{in}(X)\geq f'(X)+1=f(X)-\epsilon(X)$. If X is a $b\bar{c}$ set, then $d_D^{in}(X)\geq f'(X)\geq f(X)-\epsilon(X)$. Consider the case where X is neither a $c\bar{b}$-set or a $b\bar{c}$ set. If X is a $b\bar{a}$ set containing c, then X cannot be tight hence $d^{in}(X)\geq f'(X)+1=f(X)-\epsilon(X)$. Finally, if X is none of the above, then $d_D^{in}(X)\geq f'(X)=f(X)-\epsilon(X)$. Hence $\alpha(X)\geq 0$ and the base case is proven.

Assume that the proposition holds when $d^{out}_{P\,\mathrm{red}}(X)=n$. We prove that it also holds when $d^{out}_{P\,\mathrm{red}}(X)=n+1$. Let X be a set where $d^{out}_{P\,\mathrm{red}}(X)=n+1$ and let zy be the first red arc on P, leaving X. Note that $T(z)$ cannot contain both a and b, otherwise we could replace c by z and delete the subpath of P from z to c, contradicting the minimality of P. If $c\in T(z)$, then $T(c)\subseteq T(z)$, so $a,b\in T(z)$, and we have already noted that this is not possible. Hence we can assume that c is not in $T(z)$ and that either a or b is not in $T(z)$.

Claim 1. *Consider a set Y of vertices such that $c\notin Y$ and either $a\notin Y$ or $b\notin Y$. Then Y is tight if and only if $\alpha(Y)=0$*

If Y does not contain b, then $\epsilon(Y)=0$. On the other hand, $f(Y)=f'(Y)$. Thus $\alpha(Y)=0=d^{in}(Y)-f(Y)+\epsilon(Y)$ if and only if $d^{in}(Y)=f(Y)=f'(Y)$, that is when Y is tight.

Now, assume that Y contains b. Then Y cannot contain a. Hence $\epsilon(Y)=1$ and we have $f'(Y)=f(Y)-1$ hence Y is tight if and only if $d^{in}(Y)=f'(Y)=f(Y)-1$, that is $\alpha(Y)=0$. If Y does not contain a nor b, then $\epsilon(Y)=0$, $f(Y)=f'(Y)$, hence $\alpha(Y)=0$ if and only if $d^{in}(Y)=f(Y)$, i.e., Y is tight. The claim is proved.

Recall that $T(z)$ is the smallest tight set containing z, also recall that $T(z)$ is not contained in X because zy is a red arc leaving X, hence $T(z) \cap X$ is not tight. Note that Claim 1 applies to $T(z)$ and its subsets because $c \notin T(z)$ and $a \notin T(z)$ or $b \notin T(z)$. It follows from the claim that $\alpha(T(z)) = 0$ because $T(z)$ is tight and $\alpha(T(z) \cap X) \geq 1$, because $T(z) \cap X$ is not tight. Hence, from the submodularity of α, $\alpha(X) = \alpha(X) + \alpha(T(z)) \geq \alpha(T(z) \cap X) + \alpha(T(z) \cup X) \geq \alpha(T(z) \cup X) + 1$.

Claim 2. $d^{out}_{P\ red}(T(z) \cup X) = d^{out}_{P\ red}(X) - 1$.

For the claim, recall that no red arc leaves a tight set, hence no red arc leaves $T(z)$. Then $d^{out}_{P\ \mathrm{red}}(T(z) \cup X) \leq d^{out}_{P\ \mathrm{red}}(X) - 1$ because zy is a red arc of P leaving X and $y \in T(z)$. On the other hand, there is no other red arc vw of P leaving X with its head w in $T(z)$, otherwise we could replace the subpath $z, y..., v, w$ of P by z, w, contradicting the minimality of P; recall that zy is the first red arc of P leaving X, hence it would precede vw in P. Hence $d^{out}_{P\ \mathrm{red}}(T(z) \cup Z) = d^{out}_{P\ \mathrm{red}}(X) - 1$.

From the claim, we have that $d^{out}_{P\ \mathrm{red}}(T(z) \cup X) = d^{out}_{P\ \mathrm{red}}(X) - 1$. Then, by the induction hypothesis $d^{out}_{P\ \mathrm{red}}(T(z) \cup X) \leq \alpha(T(z) \cup X)$ and thus $\alpha(X) \geq \alpha(T(z) \cup X) + 1 \geq d^{out}_{P\ \mathrm{red}}(T(z) \cup X) + 1 = d^{out}_{P\ \mathrm{red}}(X)$.

$\square$

Chapter 3

Orienting Steiner Hypergraphs

3.1 Introduction

This chapter is based on a paper by Tamás Király and Lap Chi Lau [7]. We work with out-orientations, which we simply denote as orientations.

Given a Steiner hypergraph H with a designated set of terminal vertices $S \subseteq V(H)$ and a designated root vertex $r \in S$, we say that an orientation is *Steiner rooted k-hyperarc connected* if there exists a terminal vertex r such that there are k-hyperarc disjoint dipaths from r to every other terminal vertex. The *Steiner rooted orientation problem* is to find an orientation of a hypergraph that maximizes the connectivity between a root vertex r to every other terminal vertex. In this chapter, we study the Steiner rooted orientation problem.

Theorem 11. *([7] Theorem 1.1.) Suppose H is an undirected hypergraph, with set S of terminal vertices, and root $r \in S$. Then H has a Steiner rooted k-hyperarc-connected orientation if S is $2k$-hyperedge connected in H.*

Throughout this chapter, we use some abbreviations. We refer to a Steiner rooted k-hyperarc connected orientation, simply as *rooted k connected.* We refer to $2k$-hyperedge connected on S simply as $2k$ *connected on S*.

In Section 2 we give a brief overview of the proof. Section 3 is devoted to notation specific to this chapter. Section 4 introduces the extension property, one of the main tools used in the proof of Theorem 11, and uses it to formulate Theorem 12, a stronger version of Theorem 11. In Section 5 we study the properties of a minimal counterexample to the stronger theorem. In Section 6, we prove that the choice of the root vertex is irrelevant. Section 7 introduces the degree-specified orientation problem, a key tool for the proof of

Theorem 11. Then, in Section 8, we apply the degree-specified orientation problem to our setting. Section 9 and 10 are devoted to checking that the necessary properties of the degree-specified problem hold, hence the Steiner rooted k-hyperarc connected orientation exists.

3.2 Overview of proof

The proof of Theorem 11 contains two key ideas. One is the extension property (Definition 3.2), the other one is an extension of a theorem by Hakimi [6] on degree specified-orientations.

The extension property imposes new constraints on orientations; see Section 4 for details. To prove Theorem 11, we prove something even stronger: for a hypergraph H that is $2k$ connected on $S \subseteq V$, there exists a rooted k connected orientation that satisfies the extension property (Theorem 12). To prove this, we consider a minimal counterexample $\mathcal{H}$. By applying the extension property, we show that $\mathcal{H}$ has a nice structure; in particular the Steiner vertices of $\mathcal{H}$ are incident only to hyperedges of size 2, and moreover, no two of these vertices are adjacent.

We wish to find an orientation of $B(\mathcal{H})$ (the bipartite representation of $\mathcal{H}$) that is rooted k-connected and where hyperedge vertices have indegree one. We make a guess and assume that the Steiner vertices will have indegree equal to half of the number of incident edges. Now, the goal is to show that there exists an orientation of $B(\mathcal{H})$ that satisfies the requirement of rooted k-connectivity of terminals, and the indegree requirements of the nonterminals. For this, we first prove an extension of Hakimi's theorem on degree-specified orientations (Theorem 14 below). This gives Theorem 15 below. Using this new result, we achieve the goal by showing that there exists an orientation of $B(\mathcal{H})$ satisfying the connectivity and indegree requirements. This guarantees an orientation of $\mathcal{H}$ that is rooted k-connected, and, moreover, satisfies the extension property, contradicting the fact that $\mathcal{H}$ is a counterexample.

3.3 Notation

This section gives notation specific to this chapter.

Let $H = (V, E)$ be an arbitrary hypergraph, and $D = (V, A)$ an orientation of H. Let $\delta_D^{in}(X|\overline{Y})$ be the set of hyperarcs of H that enter X and are disjoint from Y. More precisely, $\delta_D^{in}(X|\overline{Y}) = \{a \in \delta^{in}(X) : a \cap Y = \emptyset\}$; $d_D^{in}(X|\overline{Y}) = |\delta_D^{in}(X|\overline{Y})|$ (See Figure 3.1). Note that if Y is the empty set, then $\delta_D^{in}(X|\overline{Y})$ is the same as $\delta_D^{in}(X)$. Let $\overrightarrow{E}_D(X, Y|\overline{Z})$

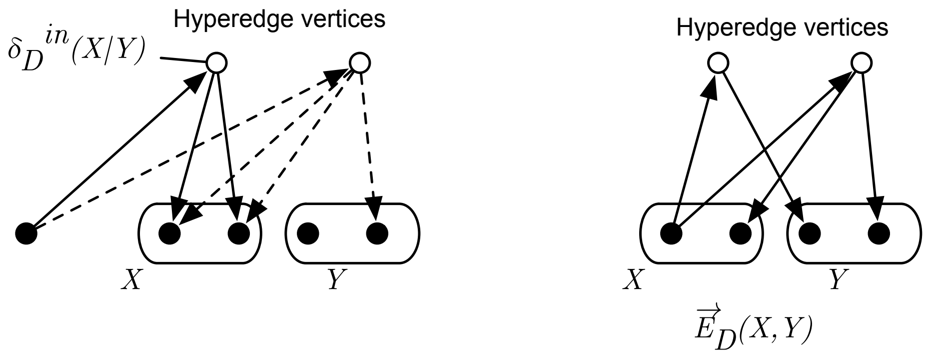

Figure 3.1: On the left, the bipartite representation of a hypergraph, a hyperedge in $\delta_D^{in}(X|\overline{Y})$ (with full line arrows) and one not in $\delta_D^{in}(X|\overline{Y})$ (with dotted arrows). On the right, the bipartite representation of another hypergraph, and the hyperedges in $\overrightarrow{E}_D(X,Y)$

denote the set of hyperarcs a of D such that a leaves X, a enters Y and $a \cap Z = \emptyset$. We use $\overrightarrow{d}_D(X,Y|\overline{Z})$ to denote $|\overrightarrow{E}(X,Y|\overline{Z})|$. If Z is the empty set, we simply use $\overrightarrow{E}_D(X,Y)$ to denote $\overrightarrow{E}_D(X,Y|\overline{Z})$ and $\overrightarrow{d}_D(X,Y)$ to denote $|\overrightarrow{E}(X,Y)|$ (See Figure 3.1).

3.4 A stronger theorem

In order to prove Theorem 11, we prove a stronger theorem, for which we need the following definition.

Definition 10. *(Figure 3.2) Given an undirected hypergraph $H = (V,E)$, $S \subseteq V$ and a vertex $s \in S$, a Steiner rooted orientation D of H* extends s *if:*

- *1. (Sink property) $d_D^{in}(s) = d_H(s)$;*
- *2. (Preflow property) $d_D^{in}(Y|\overline{s}) \geq \overrightarrow{d}_D(Y,s)$ for every $Y \subseteq V - S$.*

It is important to familiarize ourselves with this definition. The first property states that all hyperarcs of D that contain s will enter s. In other words, s will be the head of every hyperarc that contains it. The second property states that the number of hyperarcs that enter Y and do not contain s is at least the number of hyperarcs that leave Y and enter $\{s\}$. We are now ready to state the stronger theorem.

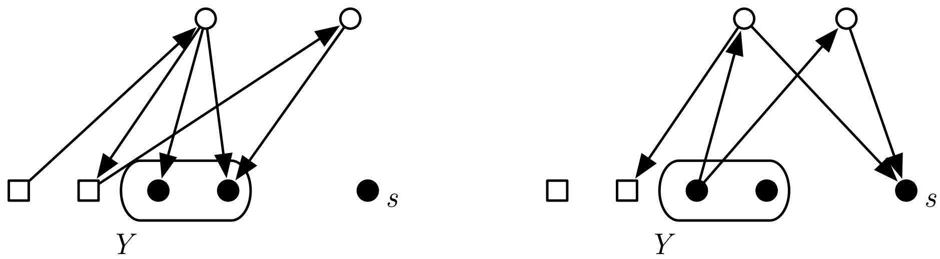

$$d_D^{in}(Y|\overline{s}) \geq \overrightarrow{d}_D(Y, s)$$

Figure 3.2: The preflow property: the number of hyperedges in $\delta_D^{in}(Y|\overline{s})$ must be at least $\overrightarrow{d}_D(Y, s)$.

Theorem 12. *Suppose H is an undirected hypergraph, S is a subset of terminal vertices with a specified root vertex $r \in S$. Then H has a Steiner rooted k-hyperarc-connected orientation if S is $2k$-hyperedge connected in H. In fact, given any vertex $s \in S$ of degree $2k$, H has a Steiner rooted k-hyperarc connected orientation that extends s.*

The only difference between Theorem 11 and Theorem 12 is that we guarantee that for a terminal vertex s, there will be a rooted k-connected orientation that extends s. This extra property is essential to the proof of Lemma 4, which is a key tool for proving Theorem 11.

3.4.1 The choice of the root is irrelevant

We use the following result stating that the choice of the root vertex is irrelevant.

Lemma 3. *Suppose there exists a Steiner rooted k-hyperarc connected orientation that extends s, with r as the root. Let v be a terminal vertex distinct from s. Then there exists a Steiner rooted k-hyperarc connected orientation that extends s with v as the root.*

We sketch a proof and defer the detailed proof to Section 3.6. Let D be a Steiner rooted k connected orientation that extends s with root r. For $v \in S - \{s, r\}$, there exist k hyperarc disjoint paths from r to v. Let D' be the orientation resulting from reversing these paths. Then it is easy to see that D' is rooted k-connected with root v. This follows because any v, r-cut has k incoming paths, and any other cut separating v from another terminal vertex, has the same number of incoming hyperarcs as in D. The extension property is also preserved, but this needs formal verification.

3.5 Properties of a minimal counterexample

Theorem 12 is proved by contradiction. Let $\mathcal{H} = (V, \mathcal{E})$ be a counterexample, with the (i) minimum number of hyperedges, and subject to (i) with the (ii) minimum number of vertices, and subject to (i) and (ii) with $\sum_{e \in \mathcal{E}} |e|$ minimum.

Definition 11. *We say that a subset X of vertices is* tight *if $d_G(X) = 2k$; we call it nontrivial if $|X| \geq 2$, $|V - X| \geq 2$, and both X and $V - X$ contain at least one terminal.*

Lemma 4. *There is no nontrivial tight set in $\mathcal{H}$.*

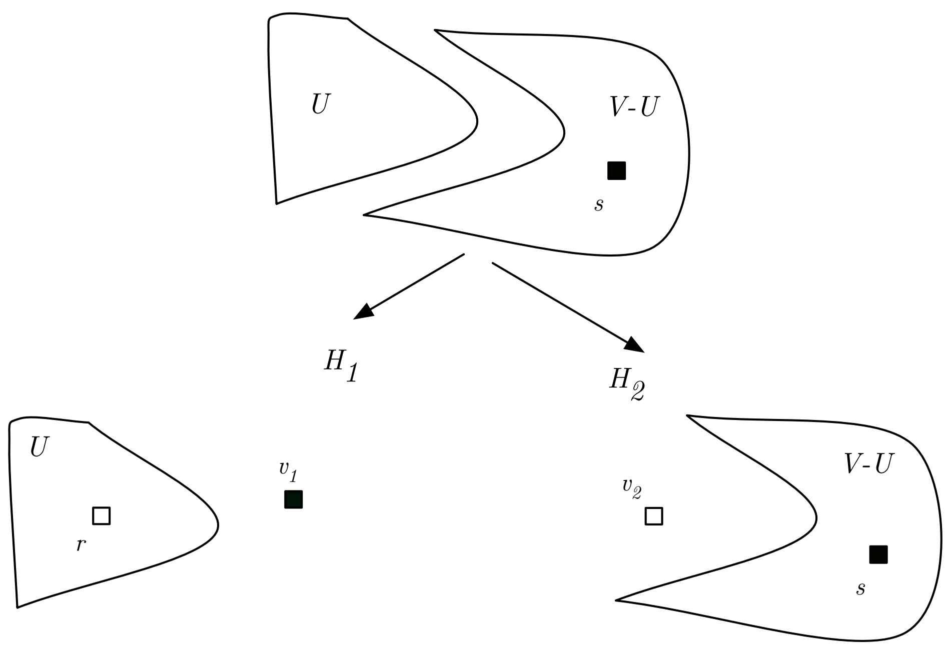

Figure 3.3: Hypergraphs H_1 and H_2 constructed from H by contracting U and $V - U$, respectively (construction for Lemma 4). The sinks of H_1 and H_2 are v_1 and s, respectively. The roots of H_1 and H_2 are r and v_2 , respectively.

We give a sketch of the proof, followed by a detailed proof. Assume U is a non trivial tight set of $\mathcal{H}$. This implies that $V - U$ is also a nontrivial tight set. For a vertex $s \in S$ of degree $2k$, we wish to find a rooted k-connected orientation of $\mathcal{H}$ that extends s. We construct two hypergraphs, one by contracting U and another by contracting $V - U$. We define the terminal vertices of these hypergraphs in a natural way. Both hypergraphs are smaller than $\mathcal{H}$, and their terminals are $2k$-connected. Since $\mathcal{H}$ is a minimal counterexample of Theorem 12, there exist rooted k-connected orientations for the smaller hypergraphs. Moreover, the extension property holds for both of the smaller hypergraphs. From these two oriented hypergraphs, we construct an orientation for $\mathcal{H}$. We then prove that the orientation is rooted k-connected and that it extends s. This orientation contradicts the minimality of $\mathcal{H}$, implying that there are no nontrivial tight sets in $\mathcal{H}$. The extension property of $\mathcal{H}$ follows from the extension property of the smaller hypergraphs. The fact that the orientation of $\mathcal{H}$ is rooted k-connected comes not only from the fact that the orientations of the smaller hypergraphs are also rooted k-connected, but also from the extension property of the smaller hypergraphs. The extension property is essential for this proof.

Proof. Suppose that U is a nontrivial tight set. Without loss of generality, we can assume that $s \notin U$ (since $V - U$ is also a nontrivial tight set), moreover, by Lemma 3 we may assume $r \in U$. Moreover, by Lemma 3, we may assume $r \in U$. Let H_1 be the hypergraph obtained from $\mathcal{H}$ by contracting $V - U$ and naming the resulting vertex v_1; and let H_2 be the hypergraph resulting from $\mathcal{H}$ by contracting U, and naming the resulting vertex v_2. Hence, $V(H_1) = U \cup \{v_1\}$ and $V(H_2) = (V - U) \cup \{v_2\}$. Each hyperedge $e \in \delta_{\mathcal{H}}(U)$ corresponds to hyperedges $e_1 = (e \cap V(H_1)) \cup \{v_1\}$ in H_1 and $e_2 = (e \cap V(H_2)) \cup \{v_2\}$ in H_2. This implies a one to one mapping between the hyperedges in $\delta_{H_1}(v_1)$ and the ones in $\delta_{H_2}(v_2)$.

Since U is nontrivial, both H_1 and H_2 are smaller than $\mathcal{H}$. Let $S_1 = (S \cap U) \cup \{v_1\}$ be the set of terminals of H_1 and $S_2 = (S \cap (V - U)) \cup \{v_2\}$ be the set of terminals of H_2. In other words, the terminals of H_1 and H_2 are the vertices coming from the terminals of $\mathcal{H}$ (both v_1 and v_2 come from at least one terminal). Note that S_1 is $2k$ connected in H_1, and so is S_2 in H_2. Moreover, v_1 has degree $2k$ in H_1, and the vertices of H_1 and H_2 have the same degrees as in $\mathcal{H}$. In particular s has degree $2k$ in H_2. Hence, from the minimality of $\mathcal{H}$, there exist rooted k-connected orientations D_1 of H_1 that extends v_1 and D_2 of H_2 that extends s. Furthermore, by Lemma 3, we can choose where each of this orientations is rooted; assume that r is the root of D_1 and v_2 is the root of D_2. Let D be an orientation of $\mathcal{H}$ given by the concatenation D_1 and D_2. By a concatenation, we mean the following:

For a hyperedge e, we define it's tail as follows:

- If $e \notin \delta_{\mathcal{H}}(U)$, e will have a corresponding hyperedge in H_1 or H_2, but not both. We define the tail of e to be the one defined for the corresponding hyperedge which is in H_1 or H_2.
- If $e \in \delta_{\mathcal{H}}(U)$, we define the tail of e to be the one defined for e_1 in D_1, where e_1 is the corresponding edge of e in H_1.

Using Menger's Theorem and the extension property for H_1 and H_2, we will verify that D is a rooted k-connected orientation that extends s.

By Menger's Theorem, D is a rooted k-connected orientation iff $d_D^{in}(X) \geq k$ for any $X \subseteq (V(\mathcal{H}))$ for which $r \notin X$ and $X \cap S \neq \emptyset$.

Let $X \subseteq V$ be such that $X \cap S \neq \emptyset$ and $r \notin X$. Assume first that $X \cap U \cap S \neq \emptyset$. Then $d_{D_1}^{in}(X - (V - U)) \geq k$ by the properties of the orientation D_1 of H_1. Since v_1 is the sink of H_1, there is no hyperarc going from $V - U$ to U in D. Hence we have $d_D^{in}(X) \geq d_{D_1}^{in}(X - (V - U)) \geq k$. Now, suppose that $X \cap U \cap S = \emptyset$. Let $X_1 = X \cap V(H_1)$ and $X_2 = X \cap V(H_2)$. If $X_1 = \emptyset$ then the properties of H_2 imply that $d_D^{in}(X) \geq k$. So we can assume that both X_1 and X_2 are non-empty, implying that $X_1 \subseteq V - S$, and X_2 has at least one terminal. Then,

$$d_D^{in}(X) \geq d_{D_1}^{in}(X_1|\overline{v_1}) + d_{D_2}^{in}(X_2) - \overrightarrow{d}_D(X_1, X_2). \tag{3.1}$$

It can be seen that the RHS is at least $d_{D_2}^{in}(X_2)$ because $\overrightarrow{d}_{D_1}(X_1, v_1) \geq \overrightarrow{d}_D(X_1, X_2)$, and by the preflow property of the extension, $d_{D_1}^{in}(X_1|\overline{v_1}) \geq \overrightarrow{d}_{D_1}(X_1, v_1) \geq \overrightarrow{d}_D(X_1, X_2)$. Hence $d_D^{in}(X) \geq d_{D_2}^{in}(X_2) \geq k$, where the second inequality holds because D_2 is rooted k-connected orientation. It is here where the preflow property of the extension is essential. Without it, the arcs counted in $\overrightarrow{d}_D(X_1, X_2)$, which do not contribute to $d_D^{in}(X)$, could be more than the arcs in $d_{D_1}^{in}(X_1|\overline{v_1})$, implying that $d_D^{in}(X) < k$.

We have proved that D is a rooted k-connected orientation of $\mathcal{H}$. It remains to check that D extends s. It is clear from our construction that s is a sink of D. Hence, we just need to verify the preflow property.

Consider a subset $Y \subseteq V - S$. Let $Y_1 = Y \cap V(H_1)$ and $Y_2 = Y \cap V(H_2)$. Then

$$d_D^{in}(Y|\overline{s}) \geq d_{D_1}^{in}(Y_1|\overline{v_1}) + d_{D_2}^{in}(Y_2|\overline{s}) - \overrightarrow{d}_D(Y_1, Y_2|\overline{s}). \tag{3.2}$$

From the preflow property of the extension property of D_1, we can verify that

$$d_{D_1}^{in}(Y_1|\overline{v_1}) \geq \overrightarrow{d}_{D_1}(Y_1, v_1) \geq \overrightarrow{d}_D(Y_1, Y_2|\overline{s}) + \overrightarrow{d}_D(Y_1, s).$$

From this and Equation 3.2, it follows that

$$d_D^{in}(Y|\overline{s}) \geq \overrightarrow{d}_D(Y_1, s) + d_{D_2}^{in}(Y_2|\overline{s})$$

By the preflow property of the extension property of D_2, we have $d_{D_2}^{in}(Y_2|\overline{s}) \geq \overrightarrow{d}_{D_2}(Y_2, s)$. Putting the last two inequalities together it follows that

$$d_D^{in}(Y|\overline{s}) \geq \overrightarrow{d}_D(Y_1, s) + \overrightarrow{d}_{D_2}(Y_2, s) = \overrightarrow{d}_D(Y, s),$$

as required. Thus D is a rooted k-connected orientation that extends s, which contradicts that $\mathcal{H}$ is a minimal counterexample of Theorem 12.

□

The extension property is essential for proving the above lemma. The sink property is used to define the concatenation of the orientations of H_1 and H_2. Without the preflow property, we would not be able to prove that the concatenation of such orientations is actually rooted k-connected.

Lemma 4 implies the following two results.

Corollary 7. *Every hyperedge of $\mathcal{H}$ of size at least 3 contains only terminal vertices.*

Proof. For a contradiction, let e be a hyperedge of size at least 3 containing the Steiner vertex v. Let $e' = e - v$, and $H' = \mathcal{H} - e + e'$. Note that H' is smaller than $\mathcal{H}$. We claim that H' cannot be $2k$-connected on S. Otherwise, there would exist an orientation of H' with the requirements of Theorem 12, which would imply an orientation of $\mathcal{H}$ with the requirements of Theorem 12.

Hence, assuming that H' is not $2k$-connected on S, it follows that there exists a set of vertices X such that both X and $V - X$ contain at least one terminal and $d_{H'}(X) < 2k$. But $d_{\mathcal{H}}(X) \geq 2k$, hence $e \in \delta_{\mathcal{H}}(X)$ and $d_{\mathcal{H}}(X) = 2k$. Then $e - v$ must be contained in either X or $V - X$ and v must be contained in the other set; then both X and $V - X$ have at least two vertices, since one contains $e - v$ and the other one contains v and a terminal vertex. This means that X is a nontrivial tight set, contradicting Lemma 4. □

The following result is proved by similar arguments.

Corollary 8. *There is no edge between two Steiner vertices in $\mathcal{H}$.*

Proof. For a contradiction, let e be an edge which connects two Steiner vertices. Then $\mathcal{H} - e$ cannot be $2k$-connected on S because that would contradict the minimality of $\mathcal{H}$ (for the same reasons as in the proof of Corollary 7). Hence, there exists a set $X \subseteq V$ such that both X and $V - X$ have at least one terminal and $d_{\mathcal{H}}(X) = 2k$, $d_{\mathcal{H}-e}(X) < 2k$, that is $e \in \delta_{\mathcal{H}}(X)$. Note that X contains a terminal vertex and one endpoint of e which is a Steiner vertex, hence $|X| \geq 2$. Analogously, $|V - X| \geq 2$. Therefore, X is a nontrivial tight set, which contradicts Lemma 4 □

3.6 Proof of Lemma 3

Recall that Lemma 3 states that, for a hypergraph H, if there exists a Steiner rooted k-connected orientation that extends s, with r as the root, then for any other terminal vertex v, distinct from r,s, there exists a Steiner rooted k-hyperarc-connected orientation that extends s with v as the root. In this section we prove this lemma.

Proof. Lemma 3.

Let D be a Steiner rooted k hyperarc connected orientation that extends s with root r. Fix $v \in S - s$. Then there exist k hyperarc disjoint paths $\{\overrightarrow{P}_1, ..., \overrightarrow{P}_k\}$ from r to v. Let D' be the orientation resulting from D by reversing $\{\overrightarrow{P}_1, ..., \overrightarrow{P}_k\}$ (defined in Section 1.9). Let $\overleftarrow{P_i}$ be the path in D' given by reversing $\overrightarrow{P}_i$. Then $\{\overleftarrow{P_1}, ..., \overleftarrow{P_k}\}$ are k hyperarc disjoint paths in D' from v to r. We will now show that D' satisfies the hypothesis of the Lemma.

First, we verify that D' is Steiner rooted k hyperarc connected. Let $X \subseteq V(H) - v$ with $X \cap S \neq \emptyset$. If $r \in X$, $\{\overleftarrow{P_1}, ..., \overleftarrow{P_k}\}$ imply that $d^{in}_{D'}(X) \geq k$. If X does not contain r, then it contains neither endpoint of each path $\overrightarrow{P_i}$. Furthermore, $\overrightarrow{P}_i$ enters and leaves X the same number of times, hence $d^{in}_{D'}(X) = d^{in}_{D}(X)$. Moreover, $d^{in}_{D} \geq k$ by assumption, hence, $d^{in}_{D'}(X) \geq k$.

Now, we verify that D' extends s. Since s is a sink in D and s is not an endpoint of any of the paths $\{\overrightarrow{P}_1, ..., \overrightarrow{P}_k\}$, hyperarcs containing s will still be incoming to s after reversing $\{\overrightarrow{P}_1, ..., \overrightarrow{P}_k\}$. Hence, s is a sink in D'. For the preflow property of the extension, consider $Y \subseteq V(H) - S$. We wish to prove that $d^{in}_{D'}(Y|\overline{s}) \geq \overrightarrow{d}_{D'}(Y, s)$. Note that $\overrightarrow{P_i}$ enters and leaves Y the same number of times, thus we may form a pairing on the hyperarcs of $\overrightarrow{P}_i$ that enter or leave Y. Let a_1 and a_2 be hyperarcs of P_i such that, in D, a_1 enters Y and a_2 leaves Y. Let D_{a_1,a_2} be the orientation given by taking the orientation of a_1 and a_2 in D', and the rest of the hyperedges oriented as in D. To check that the preflow property holds for D_{a_1,a_2} we consider four different cases, where it is easy to verify that the preflow property holds.

1. $s \in a_1$ and $s \in a_2$. Then $d^{in}_{D_{a_1,a_2}}(Y|\overline{s}) = d^{in}_{D}(Y|\overline{s}) \geq \overrightarrow{d}_D(Y, s) = \overrightarrow{d}_{D_{a_1,a_2}}(Y, s)$.
2. $s \in a_1$ and $s \notin a_2$. Then $d^{in}_{D_{a_1,a_2}}(Y|\overline{s}) = d^{in}_{D}(Y|\overline{s}) + 1 \geq \overrightarrow{d}_D(Y, s) + 1 = \overrightarrow{d}_{D_{a_1,a_2}}(Y, s)$.
3. $s \notin a_1$ and $s \in a_2$. Then $d^{in}_{D_{a_1,a_2}}(Y|\overline{s}) = d^{in}_{D}(Y|\overline{s}) - 1 \geq \overrightarrow{d}_D(Y, s) - 1 = \overrightarrow{d}_{D_{a_1,a_2}}(Y, s)$.
4. $s \notin a_1$ and $s \notin a_2$. Then $d^{in}_{D_{a_1,a_2}}(Y|\overline{s}) = d^{in}_{D}(Y|\overline{s}) \geq \overrightarrow{d}_D(Y, s) = \overrightarrow{d}_{D_{a_1,a_2}}(Y, s)$.

Recall that $\overrightarrow{P_i}$ enters and leaves Y the same number of times. By a repeated application of the above case analysis, D' has the preflow property.

$\square$

3.7 Degree-specified orientations

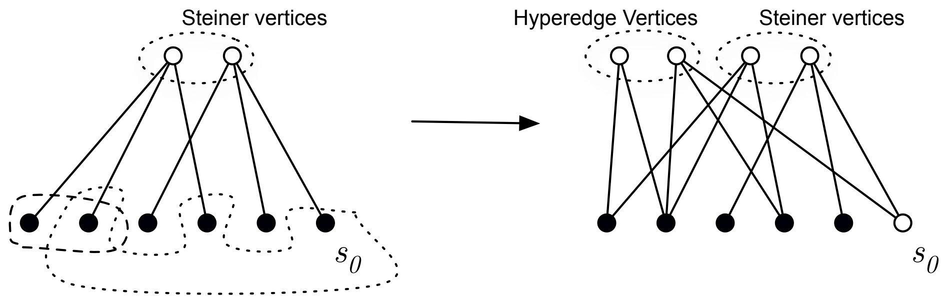

Figure 3.4: A Steiner hypergraph H on the left (as in Theorem 13). A bipartite graph B on the right, representing H. Construction done for Theorem 13. Terminal vertices are represented with full black circles.

We have proved in Corollary 7 that for $\mathcal{H}$, no Steiner vertex is incident to a hyperedge of size greater than two, that is, they are incident only to edges. Furthermore, in Corollary 8 we proved that there are no edges between Steiner vertices, thus every edge incident to a Steiner vertex has its other endpoint at a terminal. These two corollaries are key points of the proof of Theorem 12. It follows that in order to finish the proof of Theorem 12 we only need to prove following theorem.

Theorem 13. *Let $H = (V, E)$ be an undirected hypergraph. Suppose that S is 2k-connected in H, there is no edge between two Steiner vertices, and no hyperedge of size at least 3 contains a Steiner vertex. Let $s_0 \in S$ be a vertex of degree 2k. Then H has a Steiner rooted k-hyperarc-connected orientation that extends s_0.*

We will construct a bipartite representation of H. Let $B = (V', E')$ be a bipartite graph, with $V' = V \cup \mathcal{E}'$ where $\mathcal{E}'$ are the hyperedges that do not contain Steiner vertices. We say that S are the terminal vertices of B, $(V - S)$ are the Steiner vertices and $\mathcal{E}'$ the hyperedge vertices. Every terminal vertex t of B will be adjacent to the Steiner vertices adjacent to t in H and to the hyperedge vertices that contain t in $\mathcal{H}$. More specifically, $E(B) = \{ve : v \in V(H), e \in \mathcal{E}', v \in e\} \cup \{vw : v \in S, w \in V - S, vw \in \mathcal{E}\}$ (Figure 3.4).

We wish to find an orientation of B that corresponds to an orientation of H that satisfies the conclusions of Theorem 13. For this, we would need the indegree of every hyperedge vertex to be one and the indegree of s_0 to be $2k$. It is natural to assume that the Steiner vertices would have indegree $\lfloor d(v)/2 \rfloor$. It turns out that such an orientation exists. To prove it, we extend the following result by S.L Hakimi [6].

For a hypergraph $H = (V, E)$ and $X \subseteq V$, let $i(X)$ denote $|\{e \in E : e \subseteq X\}|$.

Theorem 14. *([6]) Let $G = (V, E)$ be an undirected graph, $m : V \to \mathbb{Z}$ an indegree specification, and $h : 2^V \to \mathbb{Z}^+$ a non-negative function. Then G has an orientation D such that $d_D^{in}(X) \geq h(X)$ for all $X \subseteq V$ and $d_D^{in}(v) = m(v)$ for every $v \in V$ if and only if $m(V) = |E|$ and*

$$m(X) \geq i(X) + h(X) \text{ for every } X \subseteq V.$$

We wish to extend this theorem to the setting of Steiner graphs. Unfortunately, unless NP=co-NP, this is not possible. For example, the following problem is NP-hard [7]: given a graph with a terminal set S and root r, orient the edges to maximize k, such that the orientation is Steiner rooted k-arc connected. It turns out that, if the Steiner vertices have an indegree specification, this problem is no longer NP-hard, and there exists a good characterization. We now proceed to extend Hakimi's theorem to the setting of Steiner graphs, where an indegree specification exists for the Steiner vertices, and a connectivity requirement exists for the terminal vertices. This results in Theorem 15, which is given below. Theorem 15 is used in the next section to prove Theorem 13, completing the proof of Theorem 12.

Consider a Steiner graph $G = (V, E)$ with terminal set $S \subseteq V$ and an indegree specification $m : (V - S) \to \mathbb{Z}^+$ for the Steiner vertices. For a connectivity requirement function $h : 2^S \to \mathbb{Z}^+$, we say that $h^* : 2^V \to \mathbb{Z}$ is the Steiner extension of h if $h^*(X) = h(X \cap S)$ for all $X \subseteq V$. We wish to cover h^*. From Theorem 14, a vector $m' : S \to \mathbb{Z}^+$ together with m, is the vector of indegrees of an orientation of G that covers h^* iff

- $m'(S) = |E| - m(V - S)$,
- $i(Z) \leq m(Z)$ for every $Z \subseteq V - S$
- $m'(X) + m(Z) \geq h^*(X \cup Z) + i(X \cup Z) = h(X) + i(X \cup Z)$ for every $X \subset S, Z \subseteq V - S$.

For a fixed X, the last requirement depends on all posible subsets Z of $V - S$. In other words, it is satisfied iff $m'(X) \geq h(X) + \max_{Z \subseteq V-S}(i(X \cup Z) - m(Z)))$. Let $h' : 2^S \to \mathbb{Z}^+$ be the function defined by

$$h'(X) := h(X) + \max_{Z \subseteq V-S}(i(X \cup Z) - m(Z)).$$

It follows that $m' : S \to \mathbb{Z}^+$ together with m, is the vector of indegrees of an orientation of G that covers h^* iff $m'(S) = |E| - m(V - S)$, $i(Z) \leq m(Z)$ for every $Z \subseteq V - S$ and $m'(X) \geq h'(X)$ for every $X \subset S$. If we concentrate on intersecting supermodular functions, we can determine when there exists such a vector of indegrees for the terminal vertices, guaranteeing, by Hakimi's theorem, the existence of an orientation covering h'.

Theorem 15. *Let $G = (V, E)$ be an undirected graph with a terminal set $S \subseteq V$. Let $h : 2^S \to \mathbb{Z}^+$ be a non-negative intersecting supermodular set function and $m : (V - S) \to \mathbb{Z}^+$ be an indegree specification. Then G has an orientation covering the Steiner extension h^* of h with the specified indegrees if and only if $i(Z) \le m(Z)$ for every $Z \subseteq V - S$ and for every partition $\mathcal{F}$ of S.*

$$\sum_{X \in \mathcal{F}} (h(X) + \max_{Z \subseteq V-S} (i(X \cup Z) - m(Z))) \le |E| - m(V - S).$$

We have already discussed that G has an orientation covering h^* with indegree m' at the terminal vertices iff $m'(S) = |E| - m(V - S)$, $i(Z) \le m(Z)$ for every $Z \subseteq V - S$ and $m'(X) \ge h(X) + \max_{Z \subseteq V-S}(i(X \cup Z) - m(Z)))$ for all $X \subseteq S$. It remains to show that $m' : S \to \mathbb{Z}^+$ exists if and only if

$$\sum_{X \in \mathcal{F}} (h(X) + \max_{Z \subseteq V-S} (i(X \cup Z) - m(Z))) \le |E| - m(V - S).$$

From the fact that h is intersecting supermodular, it turns out that h' is intersecting supermodular.

Lemma 5. *The set function h' is intersecting supermodular if h is intersecting supermodular.*

Proof. Let $X_1 \subseteq S$ and $X_2 \subseteq S$ be two intersecting sets. There are sets $Z_1 \subseteq V - S$ and $Z_2 \subseteq V - S$ such that $h'(X_1) = h(X_1) + i(X_1 \cup Z_1) - m(Z_1)$ and $h'(X_2) = h(X_2) + i(X_2 \cup Z_2) - m(Z_2)$. By the properties of the set functions involved, we have the following inequalities:

- $h(X_1) + h(X_2) \le h(X_1 \cap X_2) + h(X_1 \cup X_2)$.
- $i(X_1 \cup Z_1) + i(X_2 \cup Z_2) \le i((X_1 \cap X_2) \cup (Z_1 \cap Z_2)) + i((X_1 \cup X_2) \cup (Z_1 \cup Z_2))$.
- $m(Z_1) + m(Z_2) = m(Z_1 \cap Z_2) + m(Z_1 \cup Z_2)$.

Thus

$$\begin{aligned}
h'(X_1) + h'(X_2) &= h(X_1) + h(X_2) + i(X_1 \cup Z_1) + i(X_2 \cup Z_2) - m(Z_1) - m(Z_2) \\
&\le h(X_1 \cap X_2) + i(X_1 \cap X_2) \cup (Z_1 \cap Z_2) - m(Z_1 \cap Z_2) \\
&\quad + h(X_1 \cup X_2) + i(X_1 \cup X_2) \cup (Z_1 \cup Z_2) - m(Z_1 \cup Z_2) \\
&\le h'(X_1 \cap X_2) + h'(X_1 \cup X_2).
\end{aligned}$$

$\square$

Claim 3. *For a non negative intersecting supermodular set function $h : 2^S \to Z^+$, there is a vector $m' : S \to \mathbb{Z}^+$ such that $m'(X) \geq h(X)$ for all $X \subseteq S$ such that $m'(S) = |E| - m(V - S)$ if and only if*

- $h(\emptyset) = 0$,
- $\sum_{X \in \mathcal{F}} h(X) \leq |E| - m(V - S)$ *for every partition* $\mathcal{F}$ *of* S.

Using this claim with h' , Theorem 15 follows. We now prove the Claim.

Proof. The necessity is straightforward. We only need to prove the sufficiency. Let $m_0 \in \mathbb{Z}^S$ be such that $m_0(X) \geq h(X)$ for all $X \subseteq S$, and $m_0(S)$ is minimum. Let $\mathbb{F}$ be the set of all partitions of S. We will prove that $m_0(S) \leq \max_{\mathcal{F} \in \mathbb{F}} \sum_{X \in \mathcal{F}} h(X) \leq |E| - m(V - S)$. It is straightforward that if this is true, then there exists a vector $m' : S \to \mathbb{Z}^+$ such that $m' \geq m_0$ and $m'(S) = l$. Hence, it remains to show that $m_0(S) \leq \max_{\mathcal{F} \in \mathbb{F}} \sum_{X \in \mathcal{F}} h(X)$.

Let $\alpha = \max_{\mathcal{F} \in \mathbb{F}} \sum_{X \in \mathcal{F}} h(X)$. Assume that $m_0(S) > \alpha$. For every $a \in S$ we will show that there exists a set X_a such that $m_0(X_a) = h(X_a)$. Assuming the contrary, let $a \in S$ be such that X_a does not exist. Then, for every $X \subseteq S$ such that $X \ni a$, $m_0(X) > h(X)$ holds. Let $m' \in \mathbb{Z}^S$ be such that $m'(v) = m_0(v)$ for $v \in S - a$ and $m'(a) = m_0(a) - 1$. Then m' also covers h, but $m'(S) = m_0(S) - 1$, which contradicts that $m(S)$ is minimum.

Hence, for every $a \in S$, there exists a set X_a such that $m_0(X_a) = h(X_a)$. We call these subsets *tight*. Note that h is intersecting supermodular, hence, it follows that the intersection and union of intersecting tight sets is tight (see Proposition 2 in Chapter 2), and that $\cup_{a \in S} X_a = S$. This implies that there is a partition $\mathcal{P}$ of S such that every member of $\mathcal{P}$ is tight. Hence, $m_0(S) = \sum_{X \in \mathcal{P}} h(X) \leq \alpha$.

□

3.8 Using Theorem 15

We recall the bipartite representation of $H = (V, \mathcal{E})$. We say that $B = (V', E')$ is the bipartite representation of H if B is bipartite with a vertex for every vertex in V and one for every hyperedge of H that has no Steiner vertex. A partition of B is given by S and $(V - S) \cup \mathcal{E}'$ where $\mathcal{E}'$ are the hyperedges that do not contain Steiner vertices. We say that S are the terminal vertices of B, $(V - S)$ are the Steiner vertices and $\mathcal{E}'$ the hyperedge vertices. Every terminal vertex t of B will be adjacent to the Steiner vertices adjacent to t in $\mathcal{H}$ and to the hyperedge vertices that contain t in $\mathcal{H}$. More specifically, $E(B) = \{ve : v \in S, e \in \mathcal{E}', v \in e\} \cup \{vw : v \in S, w \in V - S, vw \in \mathcal{E}$ (Figure 3.5).

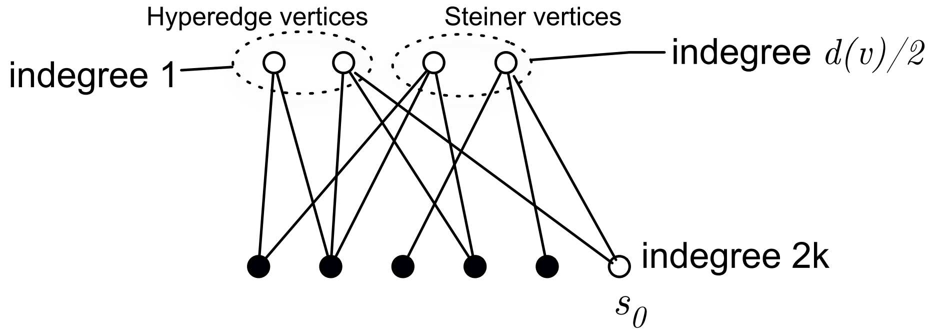

Figure 3.5: The indegree requirements for the Steiner vertices of B, the graph constructed to represent H (of Theorem 13).

We have designed Theorem 15 to work for our setting. Namely, graph B with terminal set $S' = S - s_0$ and the following indegree specification.

$$m(v) := \begin{cases} \lfloor d_H(v)/2 \rfloor & \text{if } v \text{ is a Steiner vertex} \\ 1 & \text{if } v \text{ is a hyperedge vertex} \\ 2k & \text{if } v = s_0 \text{ is the sink.} \end{cases}$$

We shall show that if B has a Steiner rooted k-connected orientation with the specified indegrees, then H has a Steiner rooted k-connected orientation that extends s_0. By Theorem 15, B has a rooted k-connected orientation with the specified indegrees if and only if the following conditions hold:

$$i(Z) \leq m(Z) \text{ for every } Z \subseteq V' - S' \tag{3.3}$$

$$\sum_{X \in \mathcal{F}} (h(X) + \max_{Y' \subseteq V' - S'} (i_B(X \cup Y) - m(Y))) \leq |E| - m(V' - S'), \tag{3.4}$$

for every partition $\mathcal{F}$ of S', where

$$h(X) = \begin{cases} k & \text{if } \emptyset \neq X \subseteq S' - r \\ 0 & \text{otherwise.} \end{cases} \tag{3.5}$$

It is easy to see that equation 3.3 is always satisfied, since the only edges spanned by $V' - S'$ are those incident to s_0 and $d_B(s_0) = 2k = m(s_0)$.

Note that $k(|\mathcal{F}|-1) = \sum_{X\in\mathcal{F}} h(X)$ for any partition $\mathcal{F}$ of S'. It follows that Equation 3.4 holds iff

$$k(|\mathcal{F}|-1) + \sum_{X\in\mathcal{F}} \max_{Y'\subseteq V'-S'} (i_B(X\cup Y) - m(Y)) \le |E| - m(V'-S'). \tag{3.6}$$

for every partition $\mathcal{F}$ of S'. Note that this inequality applies to the bipartite representation B.

Proposition 7. *Equation 3.6 holds if the following holds for the hypergraph $\mathcal{H}$. Every subpartition $\mathcal{F}$ of V such that every member of $\mathcal{F}$ has at least one terminal, and $\mathcal{F}$ restricted to S is a partition of $S - s_0$ satisfies*

$$\sum_{e\in\mathcal{E}} (|\{X\in\mathcal{F} : e\cap X \neq \emptyset| - 1) + \sum_{v\in V-(\cup_{X\in\mathcal{F}} X + s_0)} \left\lceil \frac{d_h(v)}{2} \right\rceil \ge k(|\mathcal{F}|-1). \tag{3.7}$$

Note that Equation 3.7 is in terms of $\mathcal{H}$, while Equation 3.6 was in terms of the bipartite representation B.

3.9 Proposition 7

In this section we prove Proposition 7 by contrapositive. At a glance, the equations in this section might seem complicated, but most of them come from a straight forward analysis of hypergraphs.

Proof. We will show that if Equation 3.6 does not hold, then Equation 3.7 does not hold either.

Suppose that there is a partition $\mathcal{F}$ of S' where Equation 3.6 does not hold. Hence

$$k(|\mathcal{F}|-1) + \sum_{X\in\mathcal{F}} \max_{Y\subseteq V'-S'} (i_B(X\cup Y) - m(Y))) > |E'| - m(V'-S'). \tag{3.8}$$

We will show that Equation 3.8 implies that Equation 3.6 does not hold. For $X \in \mathcal{F}$, we can determine the set Y_X where the maximum of Equation 3.8 is reached. We can assume that s_0 will not be in Y_X, since its addition to Y_X would increase $m(Y_X)$ in $2k$, while $i(X\cup Y_X)$ can increment in at most $2k$. By the same analysis, we can also assume that a Steiner vertex v will be in Y_X if and only if $|\{e\in\mathcal{E} : v\in e\subseteq X+v\}| > \lfloor d_H(v)/2\rfloor$, because v would increase $m(Y_X)$ by $\lfloor d_H(v)/2\rfloor$ and would increase $i(X\cup Y_X)$ by $|\{e\in\mathcal{E} :$

$v \in e \subseteq X + v\}|$. In the same way, a hyperedge vertex e in $\mathcal{E}'$ (recall that we defined $\mathcal{E}'$ to be the set of hyperedges that do not contain a Steiner vertex) will be in Y_X if and only if $|e \cap X| > 1$. For X, define $X^* = X \cup (Y_X \cap (V - S))$, where Y_X is the set where the maximum is attained in Equation 3.8. Note that for disjoint sets $X_1, X_2 \subseteq S - s_0$, X_1^* and X_2^* will be disjoint because a Steiner vertex v cannot have more than $\lfloor d_H(v)/2 \rfloor$ edges to two disjoint sets. Hence, $\mathcal{F}^* = \{X^* : X \in \mathcal{F}\}$ is a subpartition of V, for which each member contains a terminal, and $S \cap (\cup_{X^* \in \mathcal{F}^*} X^*) = S - s_0$. Note that no hyperedge vertex is in $(\cup_{X^* \in \mathcal{F}^*} X^*)$. Using the fact that Equation 3.8 holds for $\mathcal{F}$, then

$$k(|\mathcal{F}^*| - 1) = k(|\mathcal{F}| - 1) > |E'| - m(V' - S') - \sum_{X \in \mathcal{F}} \max_{Y \subseteq V' - S'} (i(X \cup Y) - m(Y)). \quad (3.9)$$

We will now show, that the right hand side of Equation 3.9 is at least as much as the left hand side of Equation 3.7, showing that Equation 3.7 does not hold. For this, we note some identities.

Recall that $\mathcal{E}'$ are the hyperedges of H that contain no Steiner vertex. Note that, for a hyperedge $e \in \mathcal{E}'$, there are $|e|$ edges in B, and for a hyperedge $e \in \mathcal{E} - \mathcal{E}'$ there is one edge in B. It follows that

$$|E'| = |\mathcal{E}| + \sum_{e \in \mathcal{E}'} (|e| - 1).$$

Recall that the indegree requirement m has value one on hyperedge vertices, $\lfloor \frac{d_H(v)}{2} \rfloor$ in Steiner vertices, and $2k$ in the sink s_0. Hence,

$$m(V' - S') = |\mathcal{E}'| + 2k + \sum_{v \in V - S} \left\lfloor \frac{d_H(v)}{2} \right\rfloor.$$

Recall that for a set X we defined $X^* = X \cup (Y_X \cap (V - S))$ where Y_X is the set such that $\max_{Y \subseteq V' - S'} (i_B(X \cup Y) - m(Y))$ is attained. It follows that

$$\max_{Y \subseteq V' - S'} (i_B(X \cup Y) - m(Y)) = \sum_{e \in \mathcal{E}} \max\{0, |e \cap X^*| - 1\} - \sum_{v \in X^* \cap (V - S)} \left\lfloor \frac{d_H(v)}{2} \right\rfloor.$$

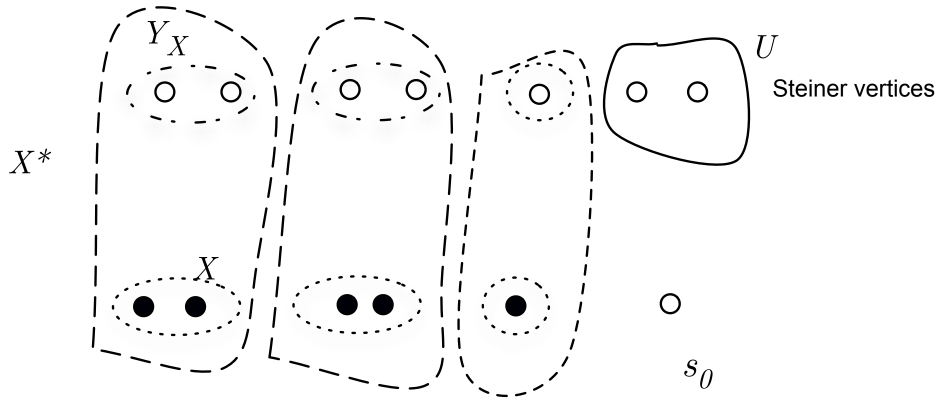

Figure 3.6: Hypergraph H, a partition $\mathcal{F}$ of the vertices $S - s_0$, and its extension to a subpartition of $V - s_0$; consider the bipartite representation B of $\mathcal{H}$; for a member X of $\mathcal{F}$, Y_X is a subset of $V' - S'$ where the maximum of Equation 3.8 is attained.

Plugging these three identities into the right hand side of Equation 3.9, we get

$$\begin{aligned} |\mathcal{E}| \quad &+ \quad \sum_{e \in \mathcal{E}'} (|e| - 1) - |\mathcal{E}'| - 2k - \sum_{v \in V - S} \left\lfloor \frac{d_H(v)}{2} \right\rfloor \\ &- \quad \sum_{X \in \mathcal{F}} \left(\sum_{e \in \mathcal{E}} \max\{0, |e \cap X^*| - 1\} - \sum_{v \in X^* \cap (V - S)} \left\lfloor \frac{d_H(v)}{2} \right\rfloor \right). \end{aligned}$$

So we started with an equation about the graph B and now we have an equation about the hypergraph $\mathcal{H}$. Let $F^* = \cup_{X^* \in \mathcal{F}^*} X^*$ and $T = (V - S) - F^*$, thus T is the set of Steiner vertices excluded from $\cup_{X^* \in \mathcal{F}^*} X^*$. Agglomerate terms together to obtain

$$|\mathcal{E}| + \sum_{e \in \mathcal{E}'} (|e| - 2) - \sum_{e \in \mathcal{E}} \sum_{X^* \in \mathcal{F}^*} \max\{0, |e \cap X^*| - 1\} - 2k - \sum_{v \in T} \left\lfloor \frac{d_H(v)}{2} \right\rfloor.$$

Note that, for a hyperedge $e \in \mathcal{E} - \mathcal{E}'$, $|e| - 1 = 1$, hence $|\mathcal{E}| + \sum_{e \in \mathcal{E}'} (|e| - 2) = \sum_{e \in \mathcal{E}} (|e| - 1)$. Hence, the last expression can be transformed into

$$\sum_{e \in \mathcal{E}} \left(|e| - (\sum_{X^* \in \mathcal{F}^*} \max\{0, |e \cap X^*| - 1\}) - 1 \right) - 2k - \sum_{v \in T} \left\lfloor \frac{d_H(v)}{2} \right\rfloor.$$

For a hyperedge $e \in \mathcal{E}$, note that $|e \cap F^*| - |\{X^* \in \mathcal{F}^* : e \cap X^* \neq \emptyset\}| = \sum_{X^* \in \mathcal{F}^*} \max\{0, |e \cap X^*| - 1\})$. Hence, we get

$$\sum_{e \in \mathcal{E}} \left(|e \cap (V - F^*)| + |\{X^* \in \mathcal{F}^* : e \cap X^* \neq \emptyset| - 1 \right) - 2k - \sum_{v \in T} \left\lfloor \frac{d_H(v)}{2} \right\rfloor.$$

Because s_0 is adjacent to $2k$ hyperedges, we can push $-2k$ into the sum as follows:

$$\sum_{e \in \mathcal{E}} \left(|e \cap (V - (F^* + s_0)| + |\{X^* \in \mathcal{F}^* : e \cap X^* \neq \emptyset\}| - 1 \right) - \sum_{v \in T} \left\lfloor \frac{d_H(v)}{2} \right\rfloor .$$

Note that $V - (F^* + s_0) = T$ and, because T consists only of Steiner vertices, every edge that intersects it does so in only one vertex, hence $\sum_{e \in \mathcal{E}} |e \cap (V - (F^* + s_0))| = \sum_{v \in T} d_H(v)$, hence $\sum_{e \in \mathcal{E}} |e \cap (V - (F^* + s_0))| - \sum_{v \in T} \left\lfloor \frac{d_H(v)}{2} \right\rfloor = \sum_{v \in T} \left\lceil \frac{d_H(v)}{2} \right\rceil$. Using this in the last equation we obtain the left hand side of Equation 3.7

$$\sum_{e \in \mathcal{F}} (|\{X^* \in \mathcal{F}^* : e \cap X^* \neq \emptyset\}| - 1) + \sum_{v \in T} \left\lceil \frac{d_H(v)}{2} \right\rceil .$$

So, we started with Equation 3.9, which states the following.

$$k(|\mathcal{F}^*| - 1) = k(|\mathcal{F}| - 1) > |E'| - m(V' - S') - \sum_{X \in \mathcal{F}} \max_{Y \subseteq V' - S'} (i(X \cup Y) - m(Y)).$$

We then proved that the right hand side of this inequality is at least the left hand side of Equation 3.7. More precisely, that it is at least

$$\sum_{e \in \mathcal{F}} (|\{X^* \in \mathcal{F}^* : e \cap X^* \neq \emptyset\}| - 1) + \sum_{v \in T} \left\lceil \frac{d_H(v)}{2} \right\rceil .$$

This shows that

$$k(|\mathcal{F}| - 1) > \sum_{e \in \mathcal{F}} (|\{X^* \in \mathcal{F}^* : e \cap X^* \neq \emptyset\}| - 1) + \sum_{v \in T} \left\lceil \frac{d_H(v)}{2} \right\rceil .$$

Hence, Equation 3.7 does not hold. □

3.10 The orientation exists if H is $2k$-connected on S.

To finish the proof of Theorem 13, we prove that Equation 3.7 holds when H is $2k$-connected on S. If Equation 3.7 holds then there exists an orientation D_B of B that is rooted k-connected with the indegree specification given by m. Let D be the orientation of $\mathcal{H}$ induced by D_B. Then D is rooted k-connected, since D_B has indegree of one at every hyperedge vertex and D_B covers the Steiner extension h^* of h. We will also check that this orientation extends s_0.

First, we prove that the bipartite representation B of H has the desired degree-specified orientation. For this, we will show that Equation 3.7 from Proposition 7 holds when H is $2k$-connected on S; recall that Equation 3.7 applies to H, not to B.

Let $\mathcal{F}$ be a subpartition of V such that every member has at least one terminal, and $\mathcal{F}$ restricted to S is a partition of $S - s_0$. Let $\mathcal{E}_1$ be the set of hyperedges from H that enter only one of the members of $\mathcal{F}$; and let $\mathcal{E}_2$ be the set of hyperedges that enter at least two parts of $\mathcal{F}$. Let $d_{\mathcal{E}_1}(X) := |\delta(X) \cap \mathcal{E}_1|$ and $d_{\mathcal{E}_2}(X) := |\delta(X) \cap \mathcal{E}_2|$. Let $U = V - (\cup_{X \in \mathcal{F}} X + s_0)$. Then the hyperedges that are disjoint from every member of $\mathcal{F}$ are the edges between U and s_0. Recall that Steiner vertices are incident only to edges that have a terminal as one of their endpoints, hence

$$\sum_{v \in U} \frac{d_H(v)}{2} = \frac{d_H(U, S)}{2} = \frac{d_H(U, S - s_0)}{2} + \frac{d_H(U, s_0)}{2} \tag{3.10}$$

For $X \in \mathcal{F}$, consider $e \in \delta(X) \cap \mathcal{E}_1$; either e is incident to s_0 or e is not incident to s_0 and so is incident to a vertex in U. Hence,

$$\sum_{X \in \mathcal{F}} d_{\mathcal{E}_1}(X) = d_H(U, S - s_0) + d_H(V - U - s_0, s_0)$$

and therefore,

$$d_H(U, S - s_0) - d_H(U, s_0) = \sum_{X \in \mathcal{F}} d_{\mathcal{E}_1}(X) - d_H(s_0) = \sum_{X \in \mathcal{F}} d_{\mathcal{E}_1}(X) - 2k \tag{3.11}$$

With the help of Equations 3.10 and 3.11, we will show, by a series of inequalities, that Equation 3.7 from Proposition 7 holds.

Recall Equation 3.7

$$\sum_{e \in \mathcal{E}} (|\{X \in \mathcal{F} : e \cap X \neq \emptyset| - 1) + \sum_{v \in V - (F + s_0)} \left\lceil \frac{d_h(v)}{2} \right\rceil \geq k(|\mathcal{F}| - 1)$$

Consider the left hand side of Equation 3.7. By taking the first sum and transforming it into a smaller sum, subtracting $d_H(U, s_0)$, and getting the ceiling operator out of the second sum we obtain the following.

$$\begin{aligned} & \sum_{e \in \mathcal{E}} (|\{X \in \mathcal{F} : e \cap X \neq \emptyset| - 1) + \sum_{v \in \cup U} \left\lceil \frac{d_h(v)}{2} \right\rceil \\ \geq \quad & \sum_{X \in \mathcal{F}} \frac{d_{\mathcal{E}_2}(X)}{2} - d_H(U, s_0) + \sum_{v \in U} \frac{d_H(v)}{2} \end{aligned}$$

By plugging in Equation 3.10 we obtain that the right hand side of the last inequality is equal to

$$\sum_{X \in \mathcal{F}} \frac{d_{\mathcal{E}_2}(X)}{2} + \frac{d_H(U, S - s_0)}{2} - \frac{d_H(U, s_0)}{2}.$$

If we plug in Equation 3.11 into the last expression, we get

$$\sum_{X \in \mathcal{F}} \left(\frac{d_{\mathcal{E}_2}(X) + d_{\mathcal{E}_1}(X)}{2} \right) - k.$$

Recall that $d_{\mathcal{E}_1}(X) + d_{\mathcal{E}_2}(X) \geq 2k$ for all $X \in \mathcal{F}$. It follows that the last expression is at least $k(|\mathcal{F}| - 1)$. This proves that Equation 3.7 from Proposition 7 holds.

From Theorem 15 and Equation 3.7, it follows that there is a rooted k-connected orientation D_B of B with the indegree specification m. This orientation induces an orientation D of H that is rooted k hyperarc connected (because every hyperedge has indegree one). To finish proving Theorem 13 we only need to check that this orientation extends s_0.

From the indegree specification, it follows that every hyperarc of D that contains s has s as a head vertex, hence s is a sink in D. It remains to see that the preflow property holds, namely that $d^{in}_{\vec{H}}(Y|\overline{s_0}) \geq \overrightarrow{d}_{\vec{H}}(Y, s_0)$ for every $Y \subseteq V - S$. So let $Y \subseteq V - S$. Since s_0 is of degree $2k$ and S is $2k$-connected in H, each vertex $v \in Y$ has at most $\lfloor d(v)/2 \rfloor$ edges to s_0 (otherwise $d(\{s_0, v\}) < 2k$ contradicting the connectivity properties). Recall that the indegree of v in the orientation is $\lfloor d(v)/2 \rfloor$ and that there are no edges between two Steiner vertices, hence all the incoming hyperarcs to a Steiner vertex $v \in Y$ come from $V - Y$. By Corollary 8 these incoming hyperarcs are of size 2, i.e., do not intersect s_0. Hence $d^{in}(Y|\overline{s_0}) \geq \overrightarrow{d}(Y, s_0)$, as needed.

Since a minimal counterexample $\mathcal{H}$ to Theorem 12 must satisfy the hypothesis of Theorem 13, then Theorem 13 proves that $\mathcal{H}$ does not exist. This proves Theorem 12 and Theorem 11.

3.11 An algorithm

Theorem 16. *Suppose $G = (V, E)$ is an undirected graph, S is a subset of terminal vertices with a specified root vertex $r \in S$, and m is an in-degree specification on the Steiner vertices (i.e. $m : (V - S) \to Z^+$). Then deciding whether G has a Steiner rooted k-arc-connected orientation with the specified indegrees can be solved in polynomial time.*

Proof. Let $h : 2^S \to \mathbb{Z}$ be the function defined by,

$$h(X) = \begin{cases} k \text{ if } \emptyset \neq X \subseteq S - r \\ 0 \text{ otherwise} \end{cases}$$

Note that h is non-negative intersecting supermodular set function. Let $h' : 2^S \to \mathbb{Z}$ be the function defined as follows.

$$h'(X) := h(X) + \max_{Z \subseteq V - S}(i(X \cup Z) - m(Z))$$

We proved in Lemma 5 that h' is non-negative intersecting supermodular. To prove the theorem, we only need to show that in polynomial time we can decide if there is a suitable indegree specification for the terminal vertices, that satisfies the hypothesis of Theorem 14. For this, we show that the the following polyhedron is box-integer, implying that, by the Complexity analysis of Theorem 60.1 [11], it can be decided in polynomial time if there is a suitable indegree specification for the terminal vertices.

$$\mathcal{B} := \{x \in \mathbb{R}^S : x(X) \geq h'(X) \text{ for } X \subseteq V, x(S) = |E| - m(V - S)\}$$

We will construct a digraph and a crossing submodular function where we will apply Edmond-Giles Theorem to show that $\mathcal{B}$ is box-integer.

Let $S' = S \cup \{z\}$, and $\mathcal{C} = 2^S \cup \{z\}$. Note that $\mathcal{C}$ is a crossing family of subsets from S'. We now define a crossing submodular function on $\mathcal{C}$.

$$f(X) = \begin{cases} -h'(X) \text{ if } X \subseteq S \\ |E| - m(V - S) \text{ if } X = \{z\} \end{cases}$$

The crossing submodularity of f comes from the intersecting supermodularity of h'. Let $D = (S', A)$, where $A = \{(z, a) : a \in S\}$. By the Edmond Giles Theorem, $\mathcal{B}' = \{x \in \mathbb{R}^A : x(\delta^{in}(U)) - x(\delta^{out}(U)) \leq f(U), \forall U \in \mathcal{C}\}$ is box integer. We will see, that if we look at $\mathbb{R}^A$ as $\mathbb{R}^S$, where the isomorphism between the two spaces is given by $(z, a) \to a$, then $\mathcal{B} = \{-x' : x' \in \mathcal{B}'\}$. Hence $\mathcal{B}$ is box integer.

Let $x \in \mathcal{B}$, define $x' = (-x)$. Note that for $U \subseteq S$, $x'(\delta^{in}(U)) - x'(\delta^{out}(U) = x'(\delta^{in}(U)) = -x(U) \leq -h'(U) = f(U)$. Also, because $x \in \mathcal{B}$, $x(S) = |E| - m(V - S)$, hence $x'(\delta^{in}(\{z\})) - x'(\delta^{out}\{z\}) = -x'(\delta^{out}\{z\}) = x(S) = |E| - m(V - S) = f(\{z\})$. Hence $x' \in \mathcal{B}'$. If $x' \in \mathcal{B}'$, then defining $x = -x'$, and the previous reasoning shows that $x \in \mathcal{B}$. □

Bibliography

[1] J. Bang-Jensen and S. Thomassé. Highly connected hypergraphs containing no two edge-disjoint spanning connected subhypergraphs. *Discrete Applied Mathematics*, 131(2):555–559, 2003. 3, 8

[2] J. Edmonds and R. Giles. A min-max relation for submodular functions on graphs. *Studies in Integer Programming. [Annals of Discrete Mathematics 1]*, pages 185–204, 1977.

[3] A. Frank. On disjoint trees and arborescences. *L. Lovász and V.T. Sós, Editors, Algebraic Methods in Graph Theory*, pages 159–169, 1981. 1, 10, 12, 13, 18

[4] A. Frank, T. Király, and Z. Király. On the orientation of graphs and hypergraphs. *Discrete Appl. Math.*, 131(2):385–400, 2003. 3, 10, 11

[5] A. Frank, T. Király, and T. Kriesell. On decomposing a hypergraph into k connected sub-hypergraphs. *Discrete Applied Mathematics*, 131(2):373–383, 2003. 9

[6] S. L. Hakimi. On the degrees of the vertices of a directed graph. *J. Franklin Inst.*, 279:290–308, 1965. 26, 34, 35

[7] T. Király and L. C. Lau. Approximate min–max theorems for steiner rooted-orientations of graphs and hypergraphs. *J. Comb. Theory Ser. B*, 98(6):1233–1252, 2008. 1, 25, 35

[8] K. Menger. Zur allgemeinen kurventheorie. *Fund. Math.*, 10:95–115, 1927. 5

[9] C. St. J. A. Nash-Williams. On orientations, connectivity and odd vertex pairings in finite graphs. *Canad. J. Math.*, 12:555–567, 1960. 8

[10] C. St. J. A. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *Journal of the London Mathematical Society*, 1(1):445, 1961. 8

[11] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency.* Springer, 2003. 6, 45

[12] W. T. Tutte. On the problem of decomposing a graph into n connected factors. *J.London Math. Soc.*, 36:221–230, 1961. 8